Spectral Clustering: An Explorative Study of Proximity Measures

Nadia Farhanaz Azam and Herna L. Viktor

School of Electrical Engineering and Computer Science, University of Ottawa
800 King Edward Avenue, Ottawa, Ontario K1N6N5, Canada
{nazam085,hlviktor}@site.uottawa.ca
http://www.eecs.uottawa.ca

Abstract. Spectral clustering algorithms recently gained much interest in research community. This surge in interest is mainly due to their ease of use, their applicability to a variety of data types and domains as well as the fact that they very often outperform traditional clustering algorithms. These algorithms consider the pair-wise similarity between data objects and construct a similarity matrix to group data into natural subsets, so that the objects located in the same cluster share many common characteristics. Objects are then allocated into clusters by employing a proximity measure, which is used to compute the similarity or distance between the data objects in the matrix. As such, an early and fundamental step in spectral cluster analysis is the selection of a proximity measure. This choice also has the highest impact on the quality and usability of the end result. However, this crucial aspect is frequently overlooked. For instance, most prior studies use the Euclidean distance measure without explicitly stating the consequences of selecting such measure. To address this issue, we perform a comparative and explorative study on the performance of various existing proximity measures when applied to spectral clustering algorithm. Our results indicate that the commonly used Euclidean distance measure is not always suitable, specifically in domains where the data is highly imbalanced and the correct clustering of boundary objects are critical. Moreover, we also noticed that for numeric data type, the relative distance measures outperformed the absolute distance measures and therefore, may boost the performance of a clustering algorithm if used. As for the datasets with mixed variables, the selection of distance measure for numeric variable again has the highest impact on the end result.

Keywords: Spectral clustering, Proximity measures, Similarity measures, Boundary detection.

1 Introduction

In cluster analysis, objects that share similar characteristics are placed into the same cluster, whereas, the objects that are very different from one another are located in two different clusters. Therefore, the main objective of a clustering algorithm is to maximize the within cluster similarity and minimize the between cluster similarity. The difference between the objects is often measured by a proximity measure, such as, similarity, dissimilarity or distance measure. There have been many clustering algorithms

A. Fred et al. (Eds.): IC3K 2011, CCIS 348, pp. 60–78, 2013.

proposed in the literature. One such family of algorithms, collectively known as the spectral clustering algorithms, recently gained much interest in the research community. This popularity is due to, amongst other, their ease of use, the fact that it can be solved efficiently by standard linear algebra software and that these algorithms very often outperforms traditional clustering algorithms such as the k-means algorithm. The algorithms from this family have been successfully deployed across numerous domains, ranging from image segmentation to clustering protein sequences. Spectral clustering algorithms have been successfully applied in a wide range of domains, ranging from oil spill detection [17] to speech separation [3].

One of the main strengths of the spectral clustering algorithm is that the algorithm may be applied to a wide range of data types (i.e. numeric, categorical, binary, and mixed) as they are not sensitive to any particular data type. These algorithms consider the pair-wise similarity between the data objects to construct a *similarity matrix*. Some of the other alternative names are *proximity*, *affinity*, or *weight matrix*. The eigenvectors and eigenvalues of this similarity matrix or the *Laplacian matrix* (a matix manipulated and constructed from the similarity matrix) are then used to find the clusters ([1], [2], [5]). The various algorithms from this family mainly differ with respect to how the similarity matrix is manipulated and/or which eigenvalue(s) and eigenvector(s) are used to partition the objects into disjoint clusters. Significant theoretical progress has been made regarding the improvement of the spectral clustering algorithms as well as the proposal of new methods, or the application in various domains. However, little research has been performed on the selection of proximity measures, which is a crucial step in constructing the similarity matrix. In this paper, we evaluate the performance of a number of such proximity measures and perform an explorative study on their behavior when applied to the spectral clustering algorithms.

Proximity measures, i.e. similarity, dissimilarity and distance measures, often play a fundamental role in cluster analysis [7]. Early steps of the majority of cluster analysis algorithms often require the selection of a proximity measure and the construction of a similarity matrix (if necessary). Most of the time, the similarity matrix is constructed from an existing similarity or distance measure, or by introducing a new measure specifically suitable for a particular domain or task. It follows that the selection of such measures, particularly when existing measures are applied, requires careful consideration as the success of these algorithms relies heavily on the choice of the proximity function ([1], [8], [9]).

Most of the previous studies on the spectral clustering algorithm use the *Euclidean distance* measure, a distance measure based on linear differences, to construct the similarity matrix for numeric feature type ([2], [5], [6]) without explicitly stating the consequences of selecting the distance measure. However, there are several different proximity measures available for numeric variable types. Each of them has their own strengths and weaknesses. To our knowledge, no in-depth evaluation of the performance of these proximity measures on spectral clustering algorithms, specifically showing that the *Euclidean distance* measure outperforms, has been carried out. As such, an evaluation and an exploratory study that compares and analyzes the performance of various proximity measures may potentially provide important guideline for researchers when selecting a proximity measure for future studies in this area. This paper endeavors to

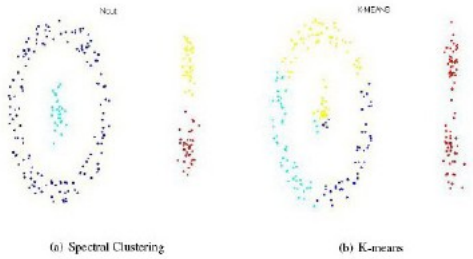

Fig. 1. Comparison of results from spectral clustering and the K-means algorithm. (a) Results from the spectral clustering algorithm, (b) Results from the K-means algorithm.

evaluate and compare the performance of these measures and to imply the conditions under which these measures may be expected to perform well.

This paper is organized as follows. In Section 2, we discuss the two spectral clustering algorithms that we used in our experiment. Section 3 presents an overview of several proximity measures for numeric, and mixed variable types. This is followed by Section 4, where we present our experimental approach and evaluate and analyze the results obtained from our experiments. We conclude the paper in Section 5.

2 Spectral Clustering

Spectral clustering algorithms originated from the area of graph partitioning and manipulate the eigenvalue(s) and eigenvector(s) of the similarity matrix to find the clusters. There are several advantages, when compared to other cluster analysis methods, to applying the spectral clustering algorithms ([1], [5], [25], [26]). Firstly, the algorithms do not make assumption on the shape of the clusters. As such, while spectral clustering algorithms may be able to find meaningful clusters with strongly coherent objects, algorithms such as K-means or K-medians may fail to do so. One such example is depicted in Figure 1 [25]. Figure 1(a) shows the ring clusters obtained from the spectral clustering algorithm and Figure 1(b) depicts the results from the K-means algorithm when applied on the same sample dataset [25]. Secondly, the algorithms do not suffer from local minima. Therefore, it may not be necessary to restart the algorithm with various initialization options. Thirdly, the algorithms are also more stable than some algorithms in terms of initializing the user-specific parameters (i.e. the number of clusters). As such, the user-specific parameters may often be estimated accurately with the help of theories related to the algorithms. Prior studies also show that the algorithms from this group thus often outperform traditional clustering algorithms, such as, K-means and Single Linkage [1]. Importantly, the algorithms from the spectral family are able to handle different types of data (i.e. numeric, nominal, binary, or mixed) and one only needs to convert the dataset into a similarity matrix to be able to apply this algorithm on a given dataset [1].

The spectral clustering algorithms are divided into two types, namely *recursive algorithms* and *multi-way algorithms* [6]. The algorithms in the first group, as the name suggest, recursively bi-partition the data at each step until a stopping criterion is satisfied. The most representative algorithm from this group is, the *Normalized Cut Spectral*

Clustering by [2]. In contrast, multi-way spectral clustering algorithms directly partition the data into k number of groups. The best-known algorithms from this group are the *Ng, Jordan and Weiss algorithm* [5] and the *Meila - Shi algorithm* [11].

In this paper, we consider two algorithms, one from each group. From the first group, we select the normalized cut spectral clustering algorithm as this algorithm proved to have had several practical successes in a variety of fields [2]. We refer to this algorithm as *SM (NCut)* in the remainder of the paper. The Ng, Jordan and Weiss algorithm is an improvement to the algorithm proposed by Meila and Shi[11] and therefore, we select this algorithm (refer to as *NJW(K-means)*) from the second group. In the following section we present several algorithm-specific notations before we discuss the algorithms themselves.

2.1 Notations

Similarity Matrix or Weight Matrix, W: Let W be an $N \times N$ symmetric, non-negative matrix where N is the number of objects in a given dataset. Let i and j be any two objects in a given dataset, located at row i and row j, respectively. If the similarity (i.e. calculated from a proximity measure) between these two objects is $w_{i,j}$, then it will be located at the cell at row i and column j in the weight matrix.

Degree Matrix, D: Let d be an $N \times 1$ matrix with $d_i = \sum_{j=1}^{n} w_{i,j}$ as the entries which denote the total similarity value from object i to the rest of the objects. Therefore, the degree matrix D is an $N \times N$ diagonal matrix which contains the elements of d on its main diagonal.

Laplacian Matrix, L**:** The Laplacian matrix is constructed from the weight matrix W and the degree matrix D. The main diagonal of this matrix is always non-negative. In graph theory, the eigenvector(s) and eigenvalue(s) of this matrix contain important information about the underlying partitions present in the graph. The spectral clustering algorithms also use the same properties to find the clusters from a given dataset.

2.2 The SM (NCut) Algorithm

The SM (NCut) spectral clustering algorithm [2], as depicted in Figure 2, is one of the most widely used recursive spectral clustering algorithm ([1], [6]). The main intuition behind this algorithm is the optimization of an objective function called the *Normalized Cut*, or *NCut*. Minimizing the NCut function is the same as finding a cut such that the total connection in between two groups is weak, whereas the total connection within each group is strong. The algorithm uses the eigenvector associated with the second smallest eigenvalue of the generalized eigenvalue system which is considered as the real valued solution to the Normalized Cut problem. The partitions are obtained by thresholding this eigenvector. There are a number of ways this grouping may be performed. One may use a particular point (i.e. zero, mean, median) as the splitting criteria or use an existing algorithm such as the K-means or K-medians algorithms for this purpose. Components with similar values usually reside in the same cluster. Since, this algorithm bi-partition

Algorithm 1: The SM (NCut) Algorithm

Input: A dataset with N number of objects.
Output: A hierarchical tree of clusters.

1. Form a $N \times N$ symmetric weight matrix W and a degree matrix D.
2. Construct the $N \times N$ Laplacian matrix $(D - W)$.
3. Solve the eigensystem $(D - W) = \lambda Dx$ for eigenvectors with smallest eigenvalues.
4. Find the second smallest eigenvalue and use the eigenvector associated with this eigenvalue to bipartition the dataset.
5. Recursively re-partition the partitions, if necessary.

Fig. 2. The algorithm for SM (NCut) spectral clustering algorithm

Algorithm 2: The NJW (K-means) Algorithm

Input: A dataset with N number of objects.
Output: A set of k number of clusters.

1. Form a $N \times N$ symmetric weight matrix W and a degree matrix D.
2. Construct the $N \times N$ Laplacian matrix $D^{\frac{-1}{2}} W D^{\frac{-1}{2}}$.
3. Find the k eigenvectors associated with k largest eigenvalues, let these eigenvectors be $x_1, x_2 \ldots x_k$.
4. Form the $N \times k$ matrix $X = [x_1 x_2 \ldots . x_k]$ by stacking the eigenvectors in columns.
5. Form another $N \times k$ matrix Y from X by normalizing each row of X so that it has a unit length. This is done by calculating $Y_{ij} = \frac{x_{ij}}{\left(\sum_j x_{ij}^2\right)^{\frac{1}{2}}}$ for each row of X.
6. Use K-means on Y to find k clusters, by treating each row of Y as a point in k-dimensions.
7. Assign the original points to the clusters formed from Y such that if a point i from Y belongs to cluster j then the original point i from the input dataset will also belong to cluster j.

Fig. 3. The algorithm for NJW(K-means) spectral clustering algorithm

the data, we get two disjoint clusters. To find more clusters we need to re-partition the segments by recursively applying the algorithm on each of the partitions.

2.3 The NJW (K-means) Algorithm

In contrast to the SM (NCut) algorithm that minimizes the *NCut* objective function and recursively bi-partitions the data, this algorithm directly partitions the data into k groups. The algorithm manipulates the normalized Laplacian matrix (as given in Figure 3) to find the clusters. The algorithm as shown in Figure 3, relates various theories from the *Random Walk Problem* and *Matrix Perturbation Theory* to theoretically motivate the partitioning solution ([1], [5], [11]). Once the eigensystem is solved and the k largest eigenvectors are normalized, the algorithm uses the *K-means* algorithm to find the k partitions.

3 Proximity Measures

Proximity measures quantify the distance or closeness between two data objects. They may be sub-categorized into three types of measures, namely *similarity*, *dissimilarity*, and *distance*.

Table 1. Proximity measures for numeric variables

Name	Function	Discussion
Euclidean Distance (EUC)	$d_{x_i,x_j} = \sqrt{\sum_{k=1}^{n}(x_{ik} - x_{jk})^2}$	Works well for compact or isolated clusters; Discovers clusters of spherical shape; Any two objects may not be influenced by the addition of a new object (i.e. outliers); Very sensitive to the scales of the variables; Not suitable for clusters of different shapes; The variables with the largest values may always dominate the distance.
Manhattan Distance (MAN)	$d_{x_i,x_j} = \sum_{k=1}^{n} \lvert x_{ik} - x_{jk} \rvert$	Computationally cheaper than the Euclidean distance; Scale dependent.
Minkowski Distance (MIN)	$d_{x_i,x_j} = \left(\sum_{k=1}^{n} \lvert x_{ik} - x_{jk} \rvert^{\lambda}\right)^{\frac{1}{\lambda}}$	One may control the amount of emphasis given on the larger differences; The Minkowski distance may cost more than the Euclidean and Manhattan distance when $\lambda > 2$.
Chebyshev Distance (CHEB)	$d_{x_i,x_j} = max_k \lvert x_{ik} - x_{jk} \rvert$	Suitable for situations where the computation time is very crucial; Very sensitive to the scale of the variables.
Canberra Distance (CAN)	$d_{x_i,x_j} = \sum_{k=1}^{n} \frac{\lvert x_{ik} - x_{jk} \rvert}{\lvert x_{ik} + x_{jk} \rvert}$	Not scale sensitive; Suitable for non-negative values; Very sensitive to the changes near the origin; Undefined when both the coordinates are 0.
Mahalanobis Distance (MAH)	$d(\boldsymbol{x}, \boldsymbol{y}) = \sqrt{(\boldsymbol{x} - \boldsymbol{y})C^{-1}(\boldsymbol{x} - \boldsymbol{y})^T}$	Considers the correlation between the variables; Not scale dependent; Favors the clusters of hyper ellipsoidal shape; Computational cost is high; May not be suitable for high-dimensional datasets.
Angular Distance (COS)	$d_{x_i,x_j} = 1 - \frac{\sum_{k=1}^{n} x_{ik} \cdot x_{jk}}{\left(\sum_{k=1}^{n} x_{ik}^2 \cdot \sum_{k=1}^{n} x_{jk}^2\right)^{\frac{1}{2}}}$	Calculates the relative distance between the objects from the origin; Suitable for semi-structured datasets (i.e. Widely applied in Text Document cluster analysis where data is highly dimensional); Does not depend on the vector length; Scale invariant; Absolute distance between the data objects is not captured.
Pearson Correlation Distance (COR)	$d_{ij} = 1 - \frac{\sum_{k=1}^{n}(x_{ik} - \bar{x_i}) \cdot (x_{jk} - \bar{x_j})}{\left(\sum_{k=1}^{n}(x_{ik} - \bar{x_i})^2 \cdot \sum_{k=1}^{n}(x_{jk} - \bar{x_j})^2\right)^{\frac{1}{2}}}$	Scale invariant; Considers the correlation between the variables; Calculates the relative distance between the objects from the mean of the data; Suitable for semi-structured data analysis (i.e. applied in microarray analysis, document cluster analysis); Outliers may affect the results.

Similarity is a numerical measure that represents the similarity (i.e. how alike the objects are) between two objects. This measure usually returns a non-negative value that falls in between 0 and 1. However, in some cases similarity may also range from -1 to $+1$. When the similarity takes a value 0, it means that there is no similarity between the objects and the objects are very different from one another. In contrast, 1 denotes complete similarity, emphasizing that the objects are identical and possess the same attribute values.

The **dissimilarity** measure is also a numerical measure, which represents the discrepancy or the difference between a pair of objects [12]. If two objects are very similar then the dissimilarity measure will have a lower value, and visa versa. Therefore, this measure is reversely related to the similarity measure. The dissimilarity value also usually fall into the interval $[0, 1]$, but it may also take values ranging from -1 to $+1$.

The term **distance**, which is also commonly used as a synonym for the dissimilarity measure, computes the distance between two data points in a multi-dimensional space. Let $d(x, y)$ be the distance between objects x and y. Then, the following four properties hold for a distance measure ([13], [14]):

1. $d(x, y) = d(y, x)$, for all points x and y.
2. $d(x, y) = 0$, if $x = y$.
3. $d(x, y) \geq 0$, for all points x and y.
4. $d(x, y) \leq d(x, z) + d(z, y)$, for all points x, y and z. This implies that introducing a third point may never shorten the distance between two other points.

There are many different proximity measures available in the literature. One of the reasons for this variety is that these measures differ on the data type of the objects in a given dataset. Next we present the proximity measures that are used in this paper.

3.1 Proximity Measures for Numeric Variables

Table 1 presents the measures for numeric, real-valued or continuous variables used in our paper ([12], [27]). These measures may be categorized into three groups. The first group contains the functions that measure the absolute distance between the objects and are scale dependent. This list includes the Euclidean (EUC), Manhattan (MAN), Minkowski (MIN), and Chebyshev (CHEB) distances. The second group contains only the Canberra distance (CAN) which also calculates the absolute distance, however, the measure is not scale dependent. In the third group we have three distance measures, namely the Angular or Cosine (COS), Pearson Correlation (COR), and Mahalanobis (MAH) distances. These measures consider the correlation between the variables into account and are scale invariant.

3.2 Proximity Measures for Mixed Variables

In the previous section, we concentrated our discussion on datasets with numeric values. Nevertheless, in practical applications, it is often possible to have more than one type of attribute in the same dataset. It follows that, in such cases, the conventional proximity measures for those data types may not work well. A more practical approach is to process all the variables of different types together and then perform a single cluster analysis [10]. The Gower's General Coefficient and Laflin's General Coefficient are two such functions that incorporate information from various data types into a single similarity coefficient. Table 2 provides the equations and additional information about these two coefficients.

4 Experiments

This section discusses our experimental methodology and the results obtained for each of the data types. In order to compare the performance of the proximity measures for a particular data type, we performed ten-fold cross validation [15] and classes to clusters evaluation on each of the datasets. In this paper, we consider the external cluster evaluation measures to measure the goodness or the quality of the clusters obtained from

Table 2. Proximity measures for mixed variables

Name	Function	Discussion
Gower's General Coefficient (GOWER)	$d(i,j) = \frac{\sum_{f=1}^{p} \delta_{ij}^{(f)} d_{ij}^{(f)}}{\sum_{f=1}^{p} \delta_{ij}^{(f)}}$	For this coefficient $\delta_{ij}^{(f)} = 0$, if one of the values is missing or the variable is of type asymmetric binary variable. Otherwise, $\delta_{ij}^{(f)} = 1$. For numeric variables distance is calculated using the formula: $d_{ij}^{(f)} = \frac{\lvert x_{if} - x_{jf} \rvert}{max_h x_{hf} - min_h x_{hf}}$. For binary and nominal variables the dissimilarity is the number of un matched pairs. This measure may be extended to incorporate other attribute types (e.g. ordinal, ratio-scaled) [14].
Laflin's General Similarity Coefficient (LAFLIN)	$s(i,j) = \frac{N_1.s_1 + N_2.s_2 + \ldots + N_n.s_n}{N_1 + N_2 + \ldots + N_n}$	In this function, $N_1 \ldots N_n$ represent the total number of attributes of each of the variable type, whereas, $s_1 \ldots s_n$ represent the total similarity measure calculated for each of the attribute type. The function uses existing similarity measures to calculate the similarity scores $s_1 \ldots s_n$.

spectral clustering algorithms, as the external class labels for each of the datasets used were available to us. Moreover, this allows us to perform a fair comparison against the known true clusters for all the proximity measures. We used two such measures namely, F-measure ([4], [16]) and G-means [17]. These measures have been previously used in numerous studies regarding cluster analysis and have proved to be successful in representing the quality of clusters numerically. In addition, we used the Friedman Test [19] to test the statistical significance of our results. The test is best suited for situations, like ours, where multiple comparisons are performed against multiple datasets. It returns the *p-value* that helps us to determine whether to accept or reject the null hypothesis. For us the null hypothesis was *"the performance of all the proximity measure is equivalent"*. For example, a p-value less than 0.05 signify that the result is only 5% likely to be extraordinary ([19], [28]).

4.1 Implementation and Settings

For all the experiments in this paper, the data preprocessing (e.g. replacing missing values, standardization) was performed using WEKA [18], an open-source Java-based machine learning software, developed at the University of Waikato in New Zealand. The spectral cluster analysis algorithms are implemented in MATLAB®. Since these algorithms manipulate the similarity matrix of a dataset, computation of eigenvalues and eigenvectors of the similarity matrix may be inefficient for a large matrix. However, MATLAB efficiently solves the eigensystem of large matrices. The cluster evaluation measures have been implemented in Java.

4.2 Datasets

Six datasets are used for each of the data types. All the datasets varied in size and are based on real-world problems representing various domains and areas. As for numeric data type, several proximity measures are scale-dependent. As such, we *standardized* [14] the attribute values to ensure the accuracy of our results. In addition, we used Equation 1 to convert the distance measure into a similarity measure in order to create the similarity or weight matrix W [2]. It is also important to note that the selection of

Table 3. Dataset Information

Numeric Datasets			
Dataset	No. of Tuples	No. of True Clusters	No. of Attributes
Body	507	2	21
Iris	150	3	4
Wine	178	3	13
Glass	214	6	8
Ecoli	336	5	7
SPECT	267	2	44
Mixed Datasets			
Dataset	No. of Tuples	No. of True Clusters	No. of Attributes
Automobile	205	6	25
CRX	690	2	15
Dermatology	366	6	33
Hepatitis	155	2	19
Post-Operative	90	3	8
Soybean	290	15	35

sigma is very crucial to the success of spectral clustering algorithm and the value of sigma varies depending on the proximity measure and the dataset used. We adopted the method proposed by Shi and Malik in [2] to select the sigma value. The sigma returned through this method was used as a starting point. Each of our experiments were performed on a range of values surrounding that sigma value and the result for which we achieved the best scores are included in this paper.

$$s(x, y) = \exp(\frac{-d(x, y)^2}{2 \times \sigma^2}) \tag{1}$$

Five of our datasets are from UCI repository [21] and one of them (Body dataset) is from an external source [20]. In Table 3, we provide a summary of the numeric datasets. The six datasets used for the mixed variable type are also obtained from the UCI repository. A summary of the mixed datasets is presented in Table 3.

4.3 Experimental Results for Numeric Datasets

In Table 4 and Table 5 we present the F-measure and G-means scores obtained for the datasets with numeric variables when SM (NCut) algorithm is used. Our results show that the COR distance and the COS distance measure often scored higher than the rest of the distance measures (Figure 4). These two distance measures performed well in four out of six datasets. The datasets are Body, Iris, Wine, and Glass. We also notice that most of the time, these two coefficients achieved similar values for both the evaluation measures. The overall average difference for these two distance measures is 0.02, irrespective of the dataset or the splitting method used. In contrast, the MAH distance measure performed poorly in four out of six datasets. The datasets for which this distance measure scored the lowest are Body, Iris, Wine, and Glass. The MAN distance performed well for the Ecoli dataset and the CAN distance performed best for

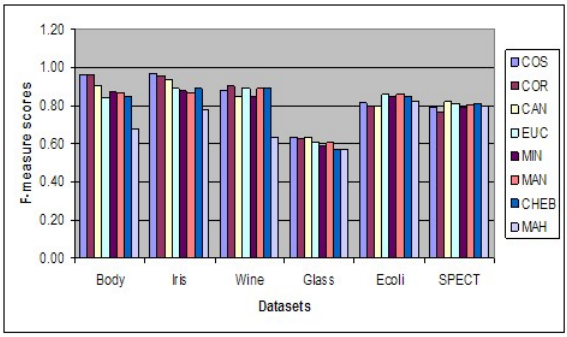

Fig. 4. F-measure scores for the numeric dataset when tested on the SM (NCut) algorithm

Table 4. F-measure scores for Numeric datasets. Algorithm: SM (NCut), Splitting points: Zero and Mean value respectively.

Dataset	COS	COR	CAN	EUC	MIN	MAN	CHEB	MAH
Body	0.97	0.96	0.90	0.87	0.87	0.87	0.85	0.68
	0.97	0.96	0.90	0.84	0.87	0.87	0.85	0.68
Iris	0.97	0.96	0.94	0.90	0.88	0.86	0.88	0.79
	0.97	0.96	0.94	0.89	0.88	0.86	0.89	0.78
Wine	0.91	0.95	0.84	0.81	0.84	0.90	0.87	0.79
	0.88	0.90	0.85	0.89	0.85	0.89	0.89	0.64
Glass	0.62	0.61	0.61	0.61	0.60	0.61	0.59	0.60
	0.64	0.63	0.64	0.61	0.59	0.61	0.57	0.57
Ecoli	0.81	0.79	0.83	0.82	0.83	0.85	0.82	0.81
	0.81	0.80	0.80	0.86	0.85	0.86	0.85	0.82
SPECT	0.80	0.77	0.81	0.80	0.80	0.81	0.80	0.80
	0.79	0.77	0.82	0.81	0.79	0.81	0.81	0.80

the SPECT dataset. We also notice that the performances of EUC, MIN, MAN, CAN and CHEB distances are very similar, and that they often scored moderately, in comparison to the highest and the lowest scores. For example, based on the scores obtained for the Body dataset, the distance measures may be grouped into three groups: 1) the COS distance and COR distance measure in one group where the scores fall in the range $[0.96 - 0.97]$, 2) the CAN, EUC, MIN, MAN and CHEB distance measures in the second group where the range is $[0.85 - 0.90]$, and 3) the MAH distance measure which scores the lowest (0.68). As observed from the results, the EUC distance measure, which is often used in the spectral cluster analysis algorithms, may not always be a suitable choice.

We observe that, if the COS distance or the COR distance measure is used instead of the EUC distance, on average the performance improved by 7.42% (F-measure) and 8.17% (G-means), respectively, for our datasets. In Table 6 and Table 7, we provide the evaluation scores from the NJW (K-means) algorithm. In this case also, the results showed almost the same trend as the results from the SM (NCut) algorithm. For both the evaluation measures, the results indicate that the MAH distance measure often scored the lowest scores over a range of datasets. Among the six datasets, in five of the cases

Table 5. G-means scores for Numeric datasets. Algorithm: SM (NCut), Splitting points: Zero and Mean value respectively.

Dataset	COS	COR	CAN	EUC	MIN	MAN	CHEB	MAH
Body	0.97	0.96	0.90	0.87	0.87	0.87	0.85	0.71
	0.97	0.97	0.90	0.83	0.84	0.87	0.85	0.71
Iris	0.97	0.96	0.94	0.90	0.88	0.87	0.88	0.80
	0.97	0.96	0.94	0.90	0.89	0.87	0.89	0.80
Wine	0.92	0.95	0.86	0.83	0.86	0.91	0.87	0.80
	0.89	0.91	0.86	0.90	0.86	0.90	0.89	0.67
Glass	0.65	0.64	0.65	0.66	0.64	0.64	0.64	0.62
	0.66	0.65	0.67	0.64	0.62	0.64	0.61	0.60
Ecoli	0.82	0.80	0.83	0.82	0.83	0.86	0.83	0.81
	0.81	0.81	0.81	0.86	0.86	0.86	0.85	0.82
SPECT	0.81	0.80	0.82	0.82	0.81	0.81	0.80	0.82
	0.81	0.80	0.83	0.82	0.82	0.82	0.81	0.81

Table 6. F-measure scores of the NJW (K-means) algorithm (tested on numeric dataset)

Dataset	COS	COR	CAN	EUC	MIN	MAN	CHEB	MAH
Body	0.88	0.88	0.83	0.79	0.79	0.80	0.78	0.64
Iris	0.81	0.82	0.79	0.78	0.80	0.78	0.79	0.69
Wine	0.94	0.96	0.97	0.85	0.97	0.86	0.83	0.50
Glass	0.54	0.54	0.53	0.55	0.56	0.58	0.54	0.50
Ecoli	0.65	0.72	0.63	0.74	0.73	0.75	0.65	0.48
SPECT	0.77	0.77	0.77	0.77	0.77	0.77	0.76	0.77

Table 7. G-means scores of the NJW (K-means) algorithm (tested on numeric dataset)

Dataset	COS	COR	CAN	EUC	MIN	MAN	CHEB	MAH
Body	0.88	0.88	0.83	0.79	0.79	0.80	0.78	0.67
Iris	0.81	0.83	0.79	0.79	0.80	0.79	0.79	0.69
Wine	0.94	0.96	0.97	0.86	0.97	0.87	0.84	0.57
Glass	0.56	0.55	0.56	0.58	0.58	0.62	0.56	0.55
Ecoli	0.69	0.74	0.68	0.75	0.75	0.77	0.68	0.52
SPECT	0.80	0.80	0.80	0.80	0.80	0.80	0.79	0.80

the MAH distance scored the lowest scores. These datasets are: Body, Iris, Wine, Glass, and Ecoli. Furthermore, in none of the cases, the EUC distance measure scored the highest score. For two datasets (i.e. Body and Iris), the COR distance and the COS distance performed well, and for the Wine and Glass datasets, the scores were very close to the highest scores achieved. The MAN distance performed well for the Ecoli and Glass datasets.

The Friedman Test which was used to measure the statistical significance of our results, gives p-values of 0.0332 (SM (NCut)) and 0.0097 (NJW (K-means)), respectively. Since the p-values are less than 0.05, this indicates the results are statistically significant. Our results showed that the MAH distance often performed poorly when compared to the rest of the distance measures according to the cluster evaluation measures. We noticed that, when the MAH distance is used, the spectral clustering algorithms produced

imbalanced clusters. Here the clusters are imbalanced when one partition contains relatively fewer objects than the other cluster. We also noticed that the objects that are placed in the smaller cluster are the objects that have the lowest degree. Recall from Section 2 that the degree is the total similarity value from one object to the rest of the objects in a dataset. In spectral clustering, the objects are considered as nodes in the graph, and a partition separates objects where the total within cluster similarity is high and the between cluster similarity is very low. Therefore, when the degree is low for an object, compared to the rest of the objects, it indicates that the object is less similar than most of the objects in the dataset. Now, the equation of the MAH distance defines an ellipsoid in n-dimensional space ([22], [23]). The distance considers the variance (how spread out the values are from the mean) of each attribute as well as the covariance (how much two variables change together) of the attributes in the datasets. It gives less weight to the dimensions with high variance and more weight to the dimensions with small variance. The covariance between the attributes allows the ellipsoid to rotate its axes and increase and decrease its size [23]. Therefore, the distance measure is very sensitive to the extreme points [24].

Figure 5 illustrates a scenario showing the MAH distance between the objects. In this figure, the distance between object 1 and 2 will be less than the distance between object 1 and 3, according to the MAH distance. This is because, object 2 lies very close to the main axes along with the other objects, whereas the object 3 lies further away from the main axes. Therefore, in such situations, the MAH distance will be large. For numeric data, the similarity will be very low when the distance is very large. The function in Equation 1, which is used to convert a distance value into a similarity value, will give a value close to zero when the distance is very large. Therefore, the degree from this object to the remainder of the objects becomes very low and the spectral methods separate these objects from the rest. This is one of the possible reasons for the MAH distance performing poorly. It either discovers imbalanced clusters or places similar objects wrongly into two different clusters. Consequently, one possible way to improve the performance of the MAH distance measure might be by changing the value of σ to a larger value. In this way, we may prevent the similarity to have a very small value.

Our results also indicate that the COR and COS distances performed best for four out of six datasets. Both of the distance measures calculate the relative distance from a fixed point (mean or zero, respectively). Therefore, two objects with a similar pattern will be more similar even if their sizes are different. The EUC, MAN, MIN, CAN and CHEB distance measures, however, calculate the absolute distance (i.e. straight line distance from one point to another). For instance, the Body dataset partitions the objects into two main clusters, one with larger body dimensions and another cluster with smaller body dimensions. According to the true cluster information, the larger body dimensions denote the Male population and the smaller body dimensions denote the Female population. When compared to the true clusters, we observed that several individuals, whose body dimensions are comparatively lower than the average body dimensions of the Male population, are placed with the individuals from the Female population by the distance measures that calculate the absolute distance. Conversely, Female individuals with larger body dimensions than the average body dimensions of the Female population are placed with the individuals from Male population. Therefore,

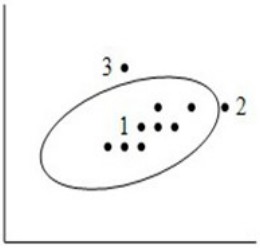

Fig. 5. A scenario depicting the Mahalanobis (MAH) distance between three points

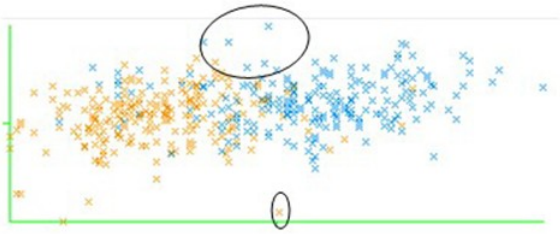

Fig. 6. Example of cluster assignments of the Body dataset. The circles are used to point to the several individual members that are placed differently.

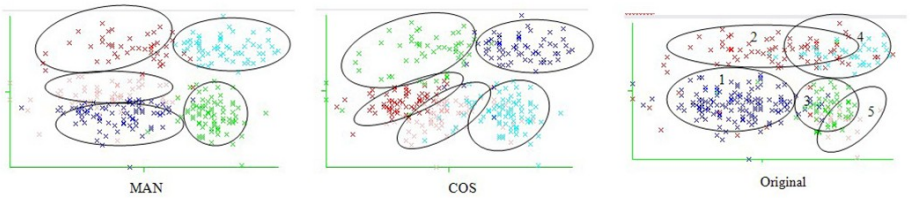

Fig. 7. Example of cluster assignments of the Ecoli dataset. (Left) The clusters obtained by using the MAN distance measure, (Middle) the clusters obtained by using the COS distance, and (Right) the original true clusters.

these individuals that fall very close to the boundary of the two true clusters, are placed differently by the distance measures that calculate the absolute distance (i.e. EUC distance, MAN distance, and MIN distance) than the distance measures that consider the relative distance (i.e. COR distance and COS distance). In such cases, the COS distance and the COR distance correctly identify these individuals. In Figure 6 we plot the first two attributes of the Body dataset when the EUC distance is used as the distance measure. The object marked with a smaller circle is an example of a Male individual with smaller body dimensions. When the EUC distance is used as the distance measure in the spectral clustering algorithm, this object is placed with the Female population. The objects marked with the larger circle illustrate the reversed situation, where Female individuals with larger body size are placed with the individuals from Male population. In both situations, the COR and COS distances placed the objects within their own groups. In Figure 7, we provide the clusters from the Ecoli dataset when the MAN (Left) and the COS (Middle) distances are used. The farthest right figure (with the title Original)

depicts the true clusters. The objects, according to the true clusters, overlap between the clusters in a number of situations (e.g. the objects marked with circle 2 and 4, or the objects marked with circle 3 and 5). This indicates that there are several objects in the dataset that may be very similar, but are placed in two different true clusters. When the spectral clustering algorithms are applied to this dataset, both the COS and MAN distances divide the true cluster marked with circle 1 (in Figure 7) into two different clusters. However, the clusters produced by the COS distance contain members from true cluster 1 and 3, whereas the clusters produced by the MAN distance contain the members from true cluster 1. The figure indicates that the shape of the clusters produced by the COS distance are more elongated toward the origin, which is the reason why some of the members from true cluster 3 are included.

Discussion. The main conclusion drawn from our results thus indicate that the MAH distance needs special consideration. This measure tends to create imbalanced clusters and therefore results in poor performance. In addition, the distance measures based on the relative distances (i.e. COR and COS distance measure) outperformed the distance measures based on the absolute distance. We noticed that, in such cases, the objects that reside in the boundary area are correctly identified by the relative distance measures. These boundary objects are slightly different from the other members of their own group and may need special attention. This is due to the fact that the COR and COS measures consider the underlying patterns in between the objects from a fixed point (i.e. mean or zero), in contrast to the absolute distance approaches. Therefore, the Euclidian (EUC) distance which is a commonly used absolute distance measure in clustering domains, may not always be a good selection for the spectral clustering algorithm, especially in domains where we are sensitive to outliers and anomalies.

In summary, under certain conditions, the GOWER similarity coefficient and the LAFLIN coefficient perform similarly. The constraints are as follows: 1) the dataset does not include asymmetric binary variables, and 2) the distance and similarity measures for each of the variables are the same. Recall from Section 4.3 that our results for numeric variables indicate that the Euclidian distance may not be the best choice for numeric variables. This choice seems to impact the performance of the LAFLIN coefficient, which may be improved by using a different distance measure for the numeric variables.

4.4 Experimental Results for Mixed Datasets

In Table 8, we present the F-measure and G-means scores for the Mixed datasets, when tests are applied on the SM (NCut) algorithm. The results from NJW (K-means) algorithm are given in Table 9. Figure 8 presents a graphical representation of our results. Our results from the external evaluation scores show that, the GOWER coefficient performed well for the Automobile and Dermatology dataset. The LAFLIN's coefficient also performed well for two of the datasets. The datasets are CRX and Hepatitis. For Post Operative and Soybean datasets, both the coefficients scored the same scores. In contrast, when the tests are applied on the NJW (K-means) algorithm, our results indicate that the GOWER coefficient performed slightly better than the LAFLIN's. In four out of six datasets the GOWER scored slightly higher scores than the LAFLIN's

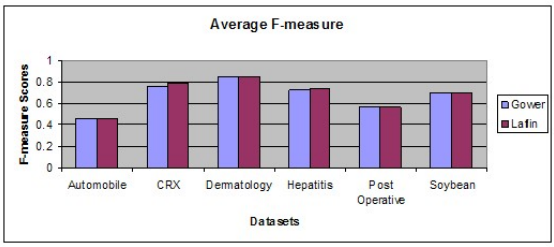

Fig. 8. Average F-measure scores for the Mixed Datasets when tested on the SM(NCut) algorithm

Table 8. F-measure and G-means scores for Mixed datasets. Algorithm: SM (NCut), Splitting points: zero and mean value, respectively.

F-measure			G-means		
Dataset	GOWER	LAFLIN	Dataset	GOWER	LAFLIN
Automobile	0.46	0.44	Automobile	0.49	0.48
	0.46	0.46		0.50	0.49
CRX	0.76	0.79	CRX	0.76	0.79
	0.76	0.79		0.76	0.79
Dermatology	0.85	0.87	Dermatology	0.86	0.87
	0.84	0.82		0.85	0.83
Hepatitis	0.71	0.73	Hepatitis	0.74	0.75
	0.73	0.75		0.75	0.77
Post Operative	0.56	0.56	Post Operative	0.57	0.57
	0.56	0.56		0.58	0.57
Soybean	0.70	0.70	Soybean	0.72	0.72
	0.70	0.70		0.72	0.72

Table 9. F-measure and G-means scores from the NJW (K-means) algorithm (tested on mixed dataset)

	F-measure		G-means	
Dataset	GOWER	LAFLIN	GOWER	LAFLIN
Automobile	0.47	0.45	0.48	0.46
CRX	0.76	0.80	0.76	0.80
Dermatology	0.84	0.82	0.86	0.84
Hepatitis	0.71	0.74	0.74	0.76
Post Operative	0.52	0.47	0.53	0.49
Soybean	0.57	0.53	0.61	0.56

coefficient. The datasets are Automobile, Dermatology, Post Operative, and Soybean. For this case also, the LAFLIN's coefficient performed best for the same two datasets (i.e. CRX and Hepatitis) as our previous test on the SM (NCut) algorithm. However, we also noticed from the scores that the difference between the performances of the two coefficients is very low. The p-values from the Friedman test are 0.3173 and 0.4142, respectively. Since, both the values are greater than 0.05, the difference between the performance of the two coefficients is not statistically significant.

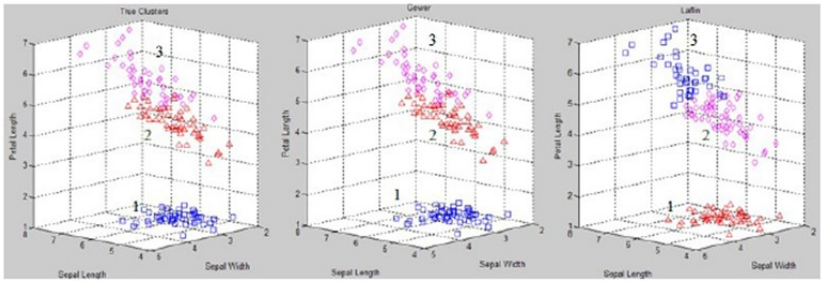

Fig. 9. Comparison of numeric functions on Iris dataset. (From left) the clusters obtained from the true clusters, the clusters obtained from the numeric function of the GOWER coefficient, and the clusters obtained from the numeric function of the LAFLIN coefficient.

In this part, we analyze the coefficients to determine the relationship between them. The equations for the two coefficients are given in Table 2. For the GOWER coefficient, the term $\delta_{ij}^{(f)}$ is an indicator variable associated with each of the variables present in the dataset and the term $d_{ij}^{(f)}$ is the distance or dissimilarity calculated for each variable for objects i and j. We also know from the description given in Table 2 that $\delta_{ij}^{(f)} = 0$ for the asymmetric binary variables and for all the other types $\delta_{ij}^{(f)} = 1$. In our datasets, all of the attributes are numeric, nominal, or symmetric binary. Therefore, the denominator of the GOWER's equation represents the total number of variables in the dataset. Also, recall that the GOWER coefficient is a dissimilarity measure, where the dissimilarity between the two objects, i and j, falls in between 0 and 1. This equation is converted into a similarity measure by subtracting from 1. Therefore, the equation for the GOWER similarity coefficient is:

$$s(i,j) = 1 - \frac{\sum_{f=1}^{p} \delta_{ij}^{(f)} d_{ij}^{(f)}}{\sum_{f=1}^{p} \delta_{ij}^{(f)}} \tag{2}$$

Let $N = \sum_{f=1}^{p} \delta_{ij}^{(f)}$ be the total number of attributes and for each attribute $\delta_{ij} = 1$, then Equation 2 becomes,

$$s(i,j) = 1 - \frac{\sum_{f=1}^{p} 1 * d_{ij}^{(f)}}{N} = \frac{N - \sum_{f=1}^{p} d_{ij}^{(f)}}{N} \tag{3}$$

Notice from the equation of LAFLIN's coefficient, s_i is the total similarity value of attribute type i, and N_i is the total number of variables of attribute type i. In our datasets, the attribute types are numeric (N_1 and s_1), nominal (N_2 and s_2), and symmetric binary (N_3 and s_3). Therefore, the equation becomes,

$$s(i,j) = \frac{N_1.s_1 + N_2.s_2 + N_3.s_3}{N_1 + N_2 + N_3} \tag{4}$$

Notice that in Equation 4, the denominator is the total number of attributes in a given dataset, which we previously denoted as N. Therefore, Equation 4 is the same as the following equation:

$$s(i,j) = \frac{N - (N_1 - N_1.s_1 + N_2 - N_2.s_2 + N_3 - N_3.s_3)}{N} \tag{5}$$

Equation 5 can be re-written as:

$$s(i,j) = \frac{N - (N_1(1 - s_1) + N_2(1 - s_2) + N_3(1 - s_3))}{N} \tag{6}$$

At this point, the GOWER equation given in Equation 3 and the LAFLIN's coefficient given in Equation 6, both have similar patterns. They have the same denominator. However, they differ only in the terms in numerator. As mentioned previously, d_{ij} is the distance or dissimilarity between the two objects i and j, whereas, $(1 - s_1)$ is also a dissimilarity measure. Both functions handle nominal and binary variables in the same way. Therefore, this implies that the difference in the equations occurs due to the functions selected for the numeric attributes which are handled differently by the two coefficients. This is one of the reasons that the difference between the performances of both of the coefficients is very low. In our tests, we used the Euclidean distance for the LAFLIN's coefficient, whereas, the GOWER coefficient uses the distance measure given in Equation 7.

$$d_{ij}^{(f)} = \frac{|x_{if} - x_{jf}|}{max_h x_{hf} - min_h x_{hf}} \tag{7}$$

We use the two numeric functions with the spectral clustering algorithms and apply them on the Iris dataset from the UCI repository [3] to evaluate their performances. Figure 9 illustrates the clusters obtained from the true clusters (left), the clusters obtained from the numeric function of the GOWER coefficient (middle), and the clusters obtained from the numeric function of the LAFLIN coefficient (right). We notice that both of the measures correctly cluster the objects from true cluster 1. However, the difference between them is clear in true cluster 2 and cluster 3. Notice that these two true clusters have objects that overlap near the boundary of the clusters. The objects located at the boundary usually have attribute values slightly different from the other members of their own true clusters. The numeric function for the GOWER coefficient correctly distinguishes several objects near the boundary. However, the LAFLIN coefficient, which used the Euclidean distance to compute the distance between the objects, placed the objects which are located near the boundary, in two different clusters. We notice that the clusters formed from this measure have a shape similar to a sphere. This may be the reason for this measure performing slightly differently than the function of the GOWER coefficient.

Discussion. In summary, under certain conditions, the GOWER similarity coefficient and the LAFLIN coefficient perform similarly. The constraints are as follows: 1) the dataset does not include asymmetric binary variables, and 2) the distance and similarity measures for each of the variables are the same. Recall from Section 4.3 that our results for numeric variables indicate that the Euclidian distance may not be the best choice for

numeric variables. This choice seems to impact the performance of the LAFLIN coefficient, which may be improved by using a different distance measure for the numeric variables.

5 Conclusions

The selection of proximity measures is a crucial step that has a significant impact on the quality and usability of the end results of the cluster analysis algorithms. This fact is frequently overlooked, leading to a degrading of the potential knowledge being discovered. This paper presented an explorative and comparative study of the performance of various proximity measures when applied to the spectral clustering algorithms. In particular, our study address the question when, and where, the choice of proximity measure becomes crucial in order to succeed. Our results indicate that proximity measures needs special care in domains where the data is highly imbalanced and where in is important to correctly cluster the boundary objects. These cases are of special interest in application areas such as rare disease diagnosis, financial market analysis and fraud detection.

Our future work will consider a diverse selection of datasets. We aim to evaluate if our conclusions hold for sparse datasets with noise and many missing values. We will also extend our research to very large datasets with high dimensionality. For such datasets, these proximity measures may not perform as per our expectation. That is, with high dimensions, the data may become sparse and the distance computed from these measures may not capture similarities properly. In such cases, a different set of proximity measures may be required to deal with the problem of high dimensionality.

The selection of the most suitable proximity measures when specifically aiming to detect outliers and anomalies is another topic of future research. In order to reach a conclusion with higher generality, we are interested to see whether the conclusions drawn from our paper persist for other clustering algorithms. The development of additional measures for mixed data types, especially ones that do not use the Euclidian distance for numeric data, are also a significant issue which will benefit from being further researched.

References

1. Luxburg, U.: A Tutorial on Spectral Clustering. Statistics and Computing 17(4), 395–416 (2007)
2. Shi, J., Malik, J.: Normalized Cuts and Image Segmentation. IEEE Transactions on Pattern Analysis and Machine Intelligence 22(8), 888–905 (2000)
3. Bach, F.R., Jordan, M.I.: Learning Spectral Clustering, with Application to Speech Separation. J. Mach. Learn. Res. 7, 1963–2001 (2006)
4. Paccanaro, A., Casbon, J.A., Saqi, M.A.: Spectral Clustering of Protein Sequences. Nucleic Acids Res. 34(5), 1571–1580 (2006)
5. Ng, A.Y., Jordan, M.I., Weiss, Y.: On Spectral Clustering: Analysis and an Algorithm. In: Dietterich, T.G., Ghahramani, S.B. (eds.) Advances in Neural Information Processing Systems, vol. 14, pp. 849–856 (2001)
6. Verma, D., Meila, M.: A Comparison of Spectral Clustering Algorithms (2001)

7. Jain, A.K., Murty, M.N., Flynn, P.J.: Data Clustering: a Review. ACM Computing Surveys 31(3), 264–323 (1999)
8. Bach, F.R., Jordan, M.I.: Learning Spectral Clustering. In: Advances in Neural Information Processing Systems 16: Proceedings of the 2003 Conference, pp. 305–312 (2003)
9. Everitt, B.S.: Cluster Analysis, 2nd edn. Edward Arnold and Halsted Press (1980)
10. Kaufman, L., Rousseeuw, P.: Finding Groups in Data: An Introduction to Cluster Analysis. Wiley-Interscience (2005)
11. Meila, M., Shi, J.: A Random Walks View of Spectral Segmentation. In: International Conference on Artificial Intelligence and Statistics (AISTAT), pp. 8–11 (2001)
12. Webb, A.R.: Statistical Pattern Recognition, 2nd edn. John Wiley & Sons (2002)
13. Larose, D.T.: Discovering Knowledge in Data: An Introduction to Data Mining. Wiley-Interscience (2004)
14. Han, J., Kamber, M.: Data Mining: Concepts and Techniques, 2nd edn. Morgan Kaufmann Publishers Inc., San Francisco (2006)
15. Costa, I.G., de Carvalho, F.A.T., de Souto, M.C.P.: Comparative Study on Proximity Indices for Cluster Analysis of Gene Expression Time Series. Journal of Intelligent and Fuzzy Systems: Applications in Engineering and Technology 13(2-4), 133–142 (2002)
16. Steinbach, M., Karypis, G., Kumar, V.: A Comparison of Document Clustering Techniques. In: KDD Workshop on Text Mining (2000)
17. Kubat, M., Holte, R.C., Matwin, S.: Machine Learning for the Detection of Oil Spills in Satellite Radar Images. Machine Learning 30(2-3), 195–215 (1998)
18. Witten, I.H., Frank, E.: Data Mining: Practical Machine Learning Tools and Techniques, 2nd edn. Morgan Kaufmann (2005)
19. Japkowicz, N., Shah, M.: Performance Evaluation for Classification A Machine Learning and Data Mining Perspective (in progress): Chapter 6: Statistical Significance Testing (2011)
20. Heinz, G., Peterson, L.J., Johnson, R.W., Kerk, C.J.: Exploring Relationships in Body Dimensions. Journal of Statistics Education 11(2) (2003)
21. Asuncion, A., Newman, D.: UCI Machine Learning Repository (2007)
22. Lee, S.-W., Verri, A. (eds.): SVM 2002. LNCS, vol. 2388. Springer, Heidelberg (2002)
23. Abou-Moustafa, K.T., Ferrie, F.P.: The Minimum Volume Ellipsoid Metric. In: Hamprecht, F.A., Schnörr, C., Jähne, B. (eds.) DAGM 2007. LNCS, vol. 4713, pp. 335–344. Springer, Heidelberg (2007)
24. Filzmoser, P., Garrett, R., Reimann, C.: Multivariate Outlier Detection in Exploration Geochemistry. Computers and Geosciences 31(5), 579–587 (2005)
25. Aiello, M., Andreozzi, F., Catanzariti, E., Isgro, F., Santoro, M.: Fast Convergence for Spectral Clustering. In: ICIAP 2007: Proceedings of the 14th International Conference on Image Analysis and Processing, pp. 641–646. IEEE Computer Society, Washington, DC (2007)
26. Fischer, I., Poland, J.: New Methods for Spectral Clustering. Technical Report IDSIA-12-04, IDSIA (2004)
27. Teknomo, K.: Similarity Measurement, `http://people.revoledu.com/kardi/tutorial/Similarity/`
28. Boslaugh, S., Watters, P.A.: Statistics in a Nutshell. O.Reilly & Associates, Inc., Sebastopol (2008)