
Clustering algorithms for intelligent web

Kanna Al Falahi, Saad Harous and
Yacine Atif*

College of Information Technology,
UAE University,
P.O. Box 17551, Al Ain, UAE
E-mail: k.alfalahi@uaeu.ac.ae
E-mail: harous@uaeu.ac.ae
E-mail: Yacine.Atif@uaeu.ac.ae
*Corresponding author

Abstract: Detecting communities of users and data in the web are important issues as the social web evolves and new information is generated everyday. In this paper, we discuss six different clustering algorithms that are related to the intelligent web. These algorithms help us to identify communities of interest in the web, which is very necessary to perform certain actions on specific group such as targeted advertisement. We consider the following algorithms: single-link algorithm, average-link algorithm, minimum-spanning-tree single-link algorithm, K-means algorithm, ROCK algorithm and DBSCAN algorithm. These algorithms are categorised into three groups: hierarchical, partitional and density-based algorithms. We show how each algorithm works and discuss potential advantages and shortcomings. We then compare these algorithms against each other and discuss their ability to accurately identify communities of interest based on social web data which are large datasets with high dimensionality. Finally, we illustrate and discuss our findings through a case study, which involves clustering in online social networks.

Keywords: algorithms; clustering; intelligent web; social networks.

Reference to this paper should be made as follows: Al Falahi, K., Harous, S. and Atif, Y. (2016) 'Clustering algorithms for intelligent web', *Int. J. Computational Complexity and Intelligent Algorithms*, Vol. 1, No. 1, pp.1–22.

Biographical notes: Kanna Al Falahi is currently a PhD candidate at the College of Information Technology (CIT) in the United Arab Emirates University. She works on analysing online social networks. She has published original contributions in web recommender systems as well as social influence maximisation in web-based social networks.

Saad Harous obtained his PhD in Computer Science from Case Western Reserve University, Cleveland, OH, USA in 1991. He has more than 20 years of experience in teaching and research in three different countries: USA, Oman and UAE. He is currently an Associate Professor at the College of Information Technology, in United Arab Emirates University. His teaching interests include programming, data structures, design and analysis of algorithms, operating systems and network. His research interests include parallel and distributed computing, and the use of computers in education and processing Arabic language. He has published more than 80 journal and conference papers. He is an IEEE senior member.

Yacine Atif received his PhD in Computer Science from Hong Kong University of Science and Technology (HKUST) in 1996. After graduation, he worked at Purdue University in the USA as a post-doc and then joined a faculty position at Nanyang Technological University (NTU) in Singapore. Since 1999 he is with the UAE University as a faculty member at the College of Information Technology. He has made a number of research contributions particularly in the areas of learning technology, semantic web and social networks. He is also involved in the technical programmes of several research forums.

1 Introduction

The web is a huge repository of information of all kinds and types. Through the years, the web has evolved dramatically. From the static Web 1.0, where webmasters create and upload web-pages with limited interaction possibilities, to the more dynamic Web 2.0, where contents are collaboratively generated and communicated across blogs, feeds and social networks. The advent of Web 3.0 brought more intelligence to web contents through the evolution of the semantic web and more automation of services over the web to support machine-to-machine interactions (SemanticWeb, 2010).

The semantic web provides novel models for retrieving and analysing web information. Using this platform, intelligent web applications are able to analyse users' inputs and behaviours to respond accordingly to different contextual considerations. These applications analyse interactions and profile users based on past history or pre-established records. Another possibility follows a case-based reasoning approach to match users with similar assets and aspirations to common web experiences. The opportunity to analyse similarities within social contexts empowers web experiences through identifying the commons to recommend preferential web contents and services (Adomavicius and Tuzhilin, 2005).

Connectivity is a core feature of intelligent web applications, where users share files, publish articles, comment on others' blogs or forums, view users' profiles and add new members to their connections. These are typical operations within today's social networks such as Facebook (<http://www.facebook.com>), MySpace (<http://www.myspace.com>) and Twitter (<http://www.twitter.com>). To make useful inferences over social connections, intelligent web applications need the following three typical modules in classical knowledge-based systems (Marmanis and Babenko, 2009):

- 1 *Content*: represented by the hypermedia data of the considered domain and composed of inter-linked resources.
- 2 *Reference*: or the knowledge base that tags and annotates domain content through rules which categorise contents into meaningful folksonomies.
- 3 *Algorithms*: which form the inference engine modules which on the domain content.

People feed the web with information everyday. This continuous flow of information may result in some inconsistencies as users will have a myriad of choices that need to be organised in an efficient manner. Data classification and clustering facilitate the process of analysing and building meaningful inferences, for example grouping similar web-pages could help finding serious problems such as mirrored web-pages or detecting

copyright violation (Haveliwala et al., 2000). In the intelligent web, there are two algorithmic approaches to categorise data: clustering and classification (Marmanis and Babenko, 2009). These approaches are useful in performing targeted advertisements or enhancing user web experience by allowing users to view posts that are of interest to them (Adomavicius and Tuzhilin, 2005) and increase individual personalised web experiences by performing special recommendations for specific users or providing page categorisations for individual users such as Google News (<http://news.google.com>).

The objective of this paper is to provide users with means to identify other users and data groups in the web. We focus on clustering algorithms in the social web context based on the following approaches categorisation: hierarchical, partitional and density-based algorithms. Each of these approaches features intrinsic techniques such as threshold or centroid techniques. We also discuss and compare six important algorithms used for clustering purpose namely: link-based (single-link, average-link and MST single-link), K-means, ROCK and DBSCAN algorithms.

The rest of this paper is organised as follows: Section 2 defines clustering. Section 3 discusses different types of clustering techniques. Section 4 introduces some important terms and concepts related to clustering. Section 5 reveals different clustering algorithms. Section 6 presents a comparison among candidate clustering algorithms. Section 7 describes a case study related to the use of clustering algorithms in social networks where we evaluate some of the candidate algorithms' clustering performance and finally Section 8 concludes this paper with a summary of works and some suggested future works.

2 Clustering in social web

Dividing people into different groups is one of human natures. Previously, people used clustering in order to study phenomena and compare them with other phenomena based on a certain set of rules. Clustering refers to grouping similar things together. It is a division of data into groups of similar objects where each group is called a cluster. It consists of objects that embody some similarities and are dissimilar to objects of other groups (Berkhin, 2002). We can find many definitions for clustering in the literatures (Jain et al., 1999; Xu and Wunsch, 2005; Gower, 1971; Jain and Dubes, 1988; Mocian, 2009; Tan et al., 2005) but the most common definition consists in partitioning data into groups (called clusters), based on some criteria so that the data grouped in one cluster should share common similarities calculated using some distance measurements. We can define clustering in the context of social network by cliques of individuals with high friendship relations internally and scattered friendship externally (Mishra et al., 2007). With clustering, we can find groups of interest that contain useful properties that are helpful to study these groups and understand their behaviours. Amazon (<http://www.amazon.com>), for example, provide users with recommendations based on their shopping experience. Twitter also started lately to recommend new friends (people to follow) to their users based on several factors, including the people these users follow and the people they follow (Discovering Who To Follow, 2010).

Clustering can be used for summarising large inputs. So instead of applying algorithms on an entire dataset, we can reduce the dataset based on specific clustering criteria (Marmanis and Babenko, 2009). Clustering analysis has been used in many research fields such as image analysis, data mining, pattern recognition, information

retrieval and machine learning (Tan et al., 2005). In the web, identifying groups of data or users would facilitate the availability and accessibility of data. Using clusters in the web is a must because of the tremendous potential which results from identifying groups of users out of a huge number of users on the internet, especially in social networks. Similarly, filtering out relevant information out of the tremendous amount of information that are linked together can lead to personalised exposure to useful information. However, in both cases, the huge size makes it hard to analyse this information.

3 Clustering types

There are many kinds of clustering algorithms available in the literatures (Jain et al., 1999; Xu and Wunsch, 2005; Mocian, 2009; Berkhin, 2002). They can be categorised based on the cluster structure (hierarchical, partitional), data types and structure (numerical, categorical) or data size (large datasets) (Marmanis and Babenko, 2009). In general, clustering approaches can be divided into four main types: hierarchical, partitional, density-based and meta-search controlled (Stein and Busch, 2005). In our paper, we discuss hierarchical, partitional and density-based clustering.

The hierarchical and partitional algorithms partition the data into different non-overlapping subsets. A partition of a dataset $X = \{x_1, x_2, \dots, x_N\}$, where $x_j = (x_{j1}, x_{j2}, \dots, x_{jd}) \in \mathfrak{R}^d$ with each measure x_{ji} called a feature (attribute, dimension or variable) and d is the input space dimensionality (Xu and Wunsch, 2005), is a collection $C = \{C_1, C_2, \dots, C_k\}$ of k non-overlapping data subsets. $C_i \neq \emptyset$ (non-null clusters) such that $C_1 \cup C_2 \cup \dots \cup C_k = X$, where X is the super cluster and $C_i \cap C_j = \emptyset$ for $i \neq j$ (Hruschka et al., 2009). The data partition is overlapping if the condition ($C_i \cap C_j = \emptyset$ for $i \neq j$) is ignored and in that case the cluster will have sub clusters of different levels inside it (Hruschka et al., 2009).

3.1 Hierarchical clustering

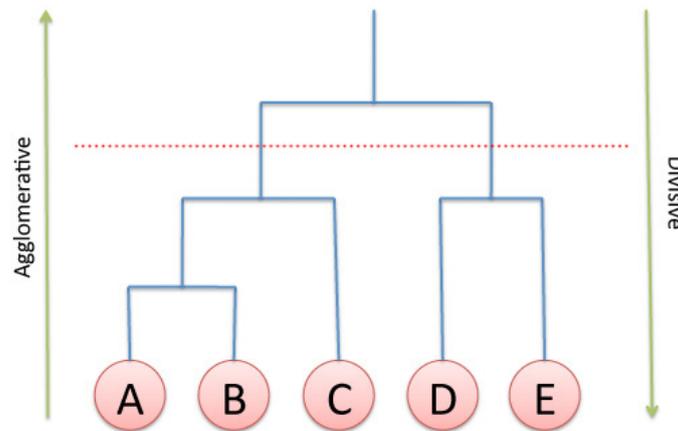
In hierarchical clustering, the clusters are represented as a tree called *dendrogram* (Xu and Wunsch, 2008). They can be either top-down (divisive) or bottom-up (agglomerative). Most of these algorithms need a threshold parameter that tells the algorithm when to stop looking for subgroups. Figure 1 shows a graphical representation of divisive and agglomerative algorithms.

In divisive hierarchical clustering, the algorithm starts from the global cluster that contains all the elements and then the data is divided into sub clusters. We need to find out which clusters to split and how to perform the splitting (Hammouda, 2002). While in agglomerative hierarchical clustering, the algorithm starts from a single cluster and then each two clusters are merged together until the global cluster is reached. DHSCAN is a hierarchical clustering algorithm represented in Naughton et al. (2006) and used to group articles that refer to the same event and have similar sentences together.

The basic idea behind clustering is to find a distance/similarity measure between any two points such as Euclidean distance, cosine distance etc. In particular, this would be the shortest path in linkage algorithms that are based on linkage metric. To calculate the distance between two points, those algorithms include single-link, average-link and MST link techniques.

Hierarchical algorithms are represented using proximity matrix (distance matrix) assuming it is symmetric which means that it require the storage of $\frac{1}{2}n^2$ proximities, where n is the number of elements (Tan et al., 2005). The total space complexity is $O(n^2)$ and the time required for computing the proximity matrix is $O(n^2)$ (Xu and Wunsch, 2008). In general, agglomerative hierarchical clustering do not have difficulties in selecting initial points as the algorithm will starts from single clusters. But they are expensive algorithms in terms of time and space which limit their usage with large-scale datasets (Xu and Wunsch, 2005). We will focus on agglomerative hierarchical algorithms in this paper such as single-link, average-link and MST single-link algorithms.

Figure 1 A dendrogram that represents divisive vs. agglomerative clustering (see online version for colours)



Note: Two clusters are generated when cutting the dendrogram at a specific level.

3.2 Partitional clustering

The partitional algorithms have fixed number of clusters where the data is divided into a number of subsets (Mocian, 2009). The most common example is the K-means algorithm that starts by selecting random means for K clusters and assign each element to its nearest mean. K-means algorithms are $O(tkn)$, where t is the number of iterations (Xu and Wunsch, 2008), k denotes the number of clusters and n the size of the data being clustered. These algorithms use a number of relocation schemes that provide optimisation to the clusters, which means the clusters can be refined at each revisiting step and thus giving an advantage over hierarchical clustering (Mocian, 2009).

3.3 Density-based clustering

In density-based algorithms, the cluster is a dense region of data objects. The points density is higher inside the cluster than outside the cluster. It is used the most when the shapes of the clusters are irregular and contain noise and outliers (Ester et al., 1996). DBSCAN is an example of density-based algorithms. In the worst case, the time complexity for this algorithm is $O(n^2)$, but in low dimensional spaces the time would be reduced to $O(n \log n)$ (Tan et al., 2005).

3.4 Meta-search controlled clustering

The meta-search controlled clustering approach treats clustering as an optimisation problem where a global goal criterion is to be minimised or maximised (Ester et al., 1996). Even though these algorithms provide flexibility, their runtime is unacceptably high. Cluster detection can be performed using genetic algorithms or two-phase greedy strategy (Ester et al., 1996).

In our paper, we will focus on hierarchical, partitional and density-based algorithms. Next, we discuss these algorithms in details.

4 Concepts and terms

4.1 Distance and similarity measures

Any clustering algorithm has a similarity factor (proximity matrix) in order to organise similar objects together. It is important to understand the measures of similarity. What makes two clusters join? What makes two points similar? And how to calculate the distance (dissimilarity)?

Xu and Wunsch (2005) defined the function of distance or dissimilarity on a dataset X in their survey paper of clustering algorithms by representing an $n * n$ symmetric proximity matrix for a dataset of n elements where the $(i, j)^{\text{th}}$ element represents the similarity or dissimilarity measure for the i^{th} and the j^{th} pattern.

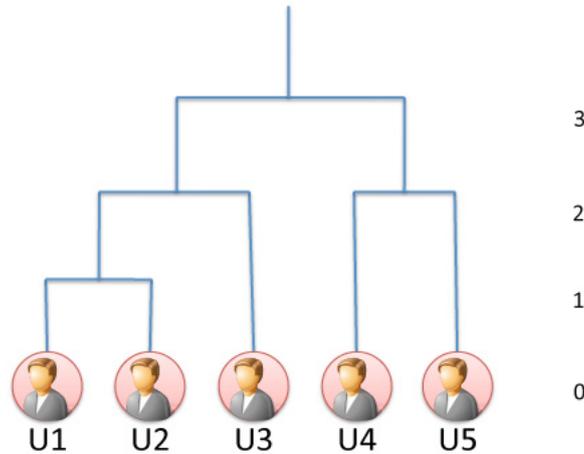
The family of Minkowski distances is a very common class of distance functions (Hammouda, 2002) and can be represented as follows:

$$D(p_i, p_j) = \sqrt[w]{\sum (p_i - p_j)} \quad (1)$$

where w is a parameter with a value greater than or equal to 1. Based on the value of w different distance functions can be represented such as Hamming distance ($w = 1$), Euclidean distance ($w = 2$) and Tschebyshev distance ($w = \infty$). Other similarity measures are *cosine correlation* measure and *Jaccard measure* (Hammouda, 2002) further discussion can be found in Xu and Wunsch (2005).

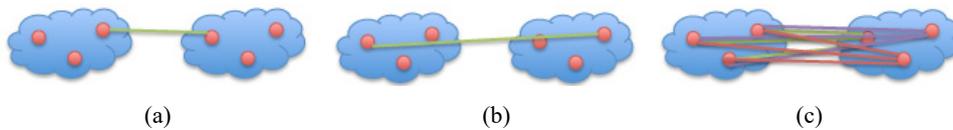
4.2 Dendrogram data structure

One of the basic structures in the clustering environment is the dendrogram, which is a tree data structure that is used to form the hierarchical cluster. Figure 2 shows a sample dendrogram with four levels. The dendrogram can be represented as a set of triples $S = \{[d, k, \{\dots\}]\}$ where d represents the threshold, k is the number of clusters and $\{\dots\}$ is the set of clusters. Figure 2 shows a dendrogram of detecting a cluster in a group of five users based on their distance similarities. The dendrogram could be represented by the following set $S = \{[0, 5, \{\{U1\}, \{U2\}, \{U3\}, \{U4\}, \{U5\}\}], [1, 4, \{\{U1, U2\}, \{U3\}, \{U4\}, \{U5\}\}], [2, 2, \{\{U1, U2, U3\}, \{U4, U5\}\}], [3, 1, \{U1, U2, U3, U4, U5\}]\}$ (Marmanis and Babenko, 2009). The dendrogram represents a set of clusters. Most of the algorithms considered in this paper are hierarchical algorithms.

Figure 2 Dendrogram structure (see online version for colours)

4.3 Proximity between clusters

Proximity calculation is the most important step in identifying clusters. It is used to measure how close are the data object to each other or how far they are, and differs according to the algorithm in use. For example, agglomerative hierarchical clustering techniques such as single-link, complete link and group average have different ways to determine the proximity threshold. The single-link defines the proximity as the closest distance between two elements in two different clusters or simply the shortest path between the two nodes in different clusters. Complete link calculates the largest distance between two points in two different clusters or the largest edge between the two nodes in different clusters. In the group average, the proximity is defined to be the average length distance of all elements from the two different clusters (Tan et al., 2005). Figure 3 illustrates the three approaches.

Figure 3 Cluster proximity, (a) single link (b) complete link (c) group average (see online version for colours)

5 An overview of clustering algorithms

In this section, we discuss and compare the following six clustering algorithms:

- 1 link-based algorithms
 - a the single-link algorithm
 - b the average-link algorithm
 - c the minimum-spanning-tree single-link algorithm

- 2 the K-means algorithm
- 3 the robust clustering using links (ROCK) algorithm
- 4 the density-based spatial clustering of applications with noise (DBSCAN) algorithm.

5.1 Link-based algorithms

The link-based algorithms are agglomerative hierarchical algorithms where the dendrogram starts with individual objects and the proximity threshold is set to zero. Then the value of threshold is increased and based on that value the algorithm checks if two elements should be merged in one cluster or be kept disjoint. After a number of iterations, all the elements will belong to a single super cluster.

The general algorithm for the hierarchical agglomerative algorithms can be described as shown in Algorithm 1.

Algorithm 1 General hierarchical agglomerative algorithm

-
- 1 Set the proximity threshold and calculate the proximity matrix
 - 2 Start with individual clusters
 - 3 Based on the threshold merge the closest clusters
 - 4 Update the threshold according to the new clusters
 - 5 Repeat steps 3 and 4 until all the elements are in one super cluster
-

The single-link, average-link and minimum-spanning-tree algorithms, which are link-based algorithms of type agglomerative hierarchical, will be discussed next.

5.1.1 The single-link

5.1.1.1 Single-link approach

The single link algorithm is based on the distance between clusters that are connected by at least one edge. First, it calculates the distance between the elements of the clusters. Then the proximity threshold is compared to the minimum distance to determine whether to merge the two clusters or not. The single-link distance between two clusters C_i and C_j could be represented by the following formula (Hammouda, 2002):

$$D(C_i, C_j) = \min_{x \in C_j, y \in C_i} (x - y) \quad (2)$$

5.1.1.2 The single-link algorithm

The single-link follows the general approach of the linked-based algorithms described in Algorithm 1. The time and space complexity of the single-link algorithm is $O(n^2)$ (Xu and Wunsch, 2005). This complexity will be a problem when working with very large data which is the case when clustering large real web datasets such as social networks.

The single link is sensitive to noise and outliers actually suffer from the chain effect (Everitt et al., 2001). This effect occurs when the single link algorithm merges two clusters based on two points in these two clusters that are close to each other, regardless of the other points of the clusters that are far away. Single-link does not provide a

solution for this problem (Marmanis and Babenko, 2009), but algorithms such as ROCK could provide a solution.

5.1.2 The average-link

5.1.2.1 Average-link approach

The average link algorithm is similar to the single-link algorithm but it uses different techniques to merge two clusters. It uses the average distance between any two points in the two different clusters and checks if it is less than the proximity threshold in order to merge the clusters.

5.1.2.2 The average-link algorithm

As with single-link algorithm, we start with individual clusters and merge them until one cluster is formed, but unlike the single link the distance of all pairs of points between the two different clusters need to be calculated.

The average distance between two clusters C_i and C_j could be represented by the following formula:

$$D(C_i, C_j) = \frac{\sum_{x \in C_j, y \in C_i} (x - y)}{C_i \cdot C_j} \quad (3)$$

The time and space complexity of the average-link algorithm is $O(n^2)$ (Xu and Wunsch, 2005). Which is similar to the single-link algorithms so it has the same problem.

5.1.3 The minimum-spanning-tree single-link

5.1.3.1 Minimum-spanning-tree single-link approach

In this approach, a minimum spanning tree connects all the elements of a given set in a way that minimises the sum of the adjacency values for the connected elements (Wu and Chao, 2004). The MST single link algorithm is a combination between the single link algorithms and the minimum spanning tree.

5.1.3.2 The minimum-spanning-tree single-link algorithm

The Prim-Jarnik algorithm (Wu and Chao, 2004) is used in this approach for the minimum spanning tree with single technique. This algorithm builds the minimum spanning tree starting from a single cluster (root) as expressed in Algorithm 2.

Algorithm 2 MST single-link algorithm

-
- 1 Mark all elements of the graph as not visited
 - 2 Choose any element you like as the root and mark it visited (cluster C created)
 - 3 The smallest-weight edge $e = (v, u)$ that connects one vertex v inside the clustering C is chosen and added to the spanning tree T.
 - 4 Repeat until all vertices are visited and the minimum spanning tree is formed
-

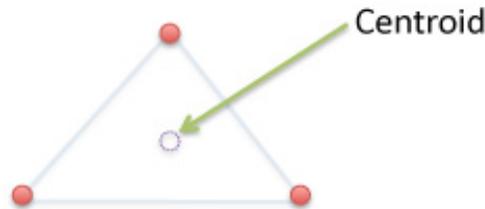
The time and space complexity of the MST single-link algorithm is $O(n^2)$ (Marmanis and Babenko, 2009). This is similar to the single-link and average-link algorithms. The MST single link algorithms results in fewer clusters than the single link algorithm because the proximity circle do not expand as much as it did in single link.

5.2 The K-means

5.2.1 K-means approach

K-means is a partitional algorithm. It uses the idea of centroid, which is the mean of a group of points. It has high performance characteristics and it is one of the oldest and most used clustering algorithms. Figure 4 illustrates the idea of centroid.

Figure 4 The centroid approach (see online version for colours)



5.2.2 The K-means algorithm

The K-means algorithm starts by choosing the K initial centroids. The simplest approach is to choose random centroids. Then the points are assigned to their closest centroid to form K clusters (Kanungo et al., 2002). Depending on the points assigned to the cluster the centroid position is updated. We repeat the update until no more points to add or the centroids remain unchanged. The K-means algorithm can be represented as shown in Algorithm 3.

Algorithm 3 K-means algorithm

-
- 1 Select K points as initial centroids
 - 2 Form K clusters by assigning each point to its closest centroid
 - 3 Re-compute the centroid of each cluster
 - 4 Repeat steps 2 and 3 until centroids are not changed
-

The K-means is fast compared to other clustering algorithms. Its computational time is $O(tkn)$ where t is the number of iterations, k represents the number of clusters and n is the number of data points we want to cluster. The space complexity for K-means is $O(n)$ which is much better than link-based algorithms.

Different runs of the K-means will produce different results since we randomly initialise the centroids actually this will produce poor clustering results. So choosing the right initial centroids is very important in order to create a good quality clusters (Kotsiantis and Pintelas, 2004). It is better to choose centroids in regions with high concentration of data points as proposed by Arthur and Vassilvitskii (2007) in their K-mean++ article.

The K-means is efficient for large datasets (Kotsiantis and Pintelas, 2004) and works well with numerical data. But the challenge occurs when it is used with categorical data such as strings since we need to find a good way to represent the nonnumeric values in a numerical way.

5.3 The ROCK

5.3.1 ROCK approach

ROCK is an agglomerative hierarchical algorithm. It uses links as similarity measure rather than measures based on distance. It clusters the points that have many common links. As an agglomerative hierarchical algorithm it starts from single clusters and merges these clusters until a super single cluster is formed. For the ROCK algorithm, we need to define a minimum number of clusters that we want to form in order to stop the algorithm before all the elements are grouped in one single cluster.

5.3.2 Goodness measure

In the process of merging clusters in the ROCK algorithm, we need to determine the best pair of clusters to merge together. Thus, the goodness measure is used. Actually for ROCK algorithm, the best clusters are those that maximise the goodness measure. The goodness measure for two clusters C_i and C_j is represented as follows (Rajeev et al., 1999):

$$g(C_i, C_j) = \frac{\text{link}[C_i, C_j]}{(n_i + n_j)^{1+2f(\theta)} - n_i^{1+2f(\theta)} - n_j^{1+2f(\theta)}} \quad (4)$$

where $\text{link}[C_i, C_j]$ represents the number of links between clusters C_i and C_j that is

$$\sum \text{link}(p_a, p_r) \quad (5)$$

Any pairs of clusters that will maximise the goodness measure will be the best pairs to merge. With algorithms that are based on similarity distance only, it will be difficult to determine if two clusters are separate because this kind of measurement may merge two clusters if there are two points close together even though these points do not have large number of common neighbours (Xu and Wunsch, 2008). Thus, the ROCK algorithm uses links as its name implies. There is a link between two data points if a common neighbour exists between them. For the ROCK algorithm to merge two clusters the focus will be on the number of links n_i, n_j between all pairs points of the two clusters C_i, C_j . The large number of links should indicate a higher probability that the two points belong to the same cluster and should give the best cluster.

5.3.3 The ROCK algorithm

The ROCK algorithm needs the following arguments:

- 1 the set of points that we want to cluster
- 2 the minimum number of clusters to have to stop the ROCK algorithm before all points are merged in one cluster

3 proximity that is required between two points in order to form a link between them.

The ROCK algorithm could be expressed as shown in Algorithm 4.

Algorithm 4 ROCK algorithm

-
- 1 Create a cluster for each point
 - 2 Use goodness measure to evaluate if two clusters should be merged or not (the best are the one those maximise the value of goodness measure)
 - 3 Repeat step 2 until the number of clusters formed is equal to the minimum number to stop or the number of cluster doesn't change between iterations
-

The space complexity of the ROCK algorithm is $O(n^2)$ (Marmanis and Babenko, 2009). While the time complexity is $O(n^2 \log n)$ (Rajeev et al., 1999).

ROCK algorithm is best used with categorical data such as keywords, Boolean attributes it uses the Jaccard coefficient to measure the similarity (Kotsiantis and Pintelas, 2004) and it works well on large dataset. One of the advantages of using the ROCK algorithm is its ability to handle outliers effectively. Outliers are point that lies in a far distance from the other points. Which means these points can be easily discard, as they will never participate in the clustering process (Rajeev et al., 1999).

5.4 The DBSCAN

5.4.1 DBSCAN approach

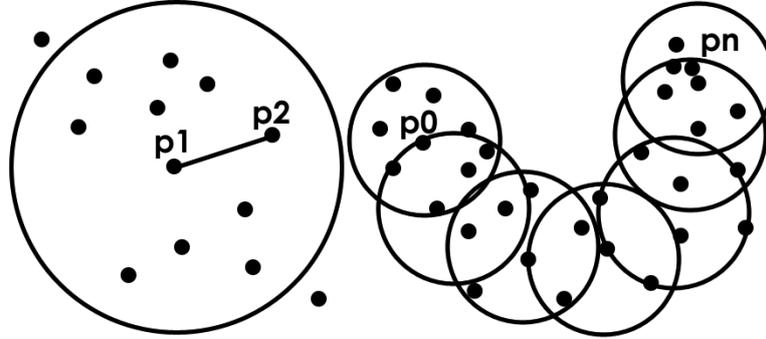
The DBSCAN algorithm is a density-based algorithm that uses density as a measurement other than links or distance between points.

Density-based algorithms are based on using density to identify the boundaries of objects. So clusters are identified based on points density within a specific region. Figure 5 explains this concept where we can identify three clusters in the figure. The points that do not belong to the clusters are identified as noise and DBSCAN is used to discover clusters and noise in a dataset.

Figure 5 Density-based clustering (see online version for colours)



The DBSCAN can be described as follows (Figure 6): any two core points should be put in the same cluster if they are close to each other within a distance Eps (Xu and Wunsch, 2008). Where Eps , stands for epsilon, is a value that helps to define an epsilon neighbourhood for any given data point p (Ester et al., 1996).

Figure 6 DBSCAN core points, border points and noise

To understand the concept of centre points let us check Figure 6 that is similar to the figure in Ester et al. (1996). The large circles are the epsilon neighbourhood for points p and q each of them is a centre of one of the circles. The circle radius is Eps and $minPoints$ represents the minimum number of points that must be inside the circle for a data point to be considered a core point. The points that are on the border of the cluster are called border points. A point p_1 is directly density-reachable from a data point p_2 with respect to Eps and $minPoints$ if there is a list of points $p_1, \dots, p_n, p_1 = q, p_n = p$ such that p_{i+1} is directly density-reachable from p_i (Ester et al., 1996). The following two conditions should be met:

- 1 p_1 inside the epsilon neighbour of p_2
- 2 there are more than $minPoints$ data points inside the epsilon neighbourhood of p_2 .

5.4.2 The DBSCAN algorithm

The DBSCAN algorithm is expressed in Algorithm 5.

Algorithm 5 DBSCAN algorithm

-
- 1 Define the points as core, border, or noise points.
 - 2 Eliminate noise points.
 - 3 Put an edge between all core points that are within Eps of each other.
 - 4 Make each group of connected core points into a separate cluster.
 - 5 Assign each border point to one of the clusters of its associated core points.
-

The time complexity for the DBSCAN algorithm is $O(n^2)$ (DBSCAN, 2012), where n is the number of points. DBSCAN can handle noise and different shape clusters because it is based on density. It can discover many clusters that are not found by the K-means algorithm. But this algorithm will have problems with clusters of very different densities as the algorithm requires that the object neighbours have enough high density (Xu and Wunsch, 2005) and with high-dimensional data (Tan et al., 2005). The DBSCAN uses R*-tree in order to improve the queries of determining the points within the Eps distance (Kotsiantis and Pintelas, 2004). R*-tree will reduce the time complexity of the DBSCAN to $O(n \log n)$ (Ester et al., 1996).

6 Analysis and discussion

In this section, we discuss the complexity of clustering algorithms and other related issues. There are many criteria that decide the use of one algorithm over the others. The two main criteria are time complexity of these algorithms and do they handle data with high dimensionality.

6.1 Large datasets

To deal with a large number of elements we need to evaluate the computational complexity of the algorithm under consideration, in other words how long does this algorithm take to construct the cluster. There is a big difference between clustering groups of people on Facebook with millions of registered users and clustering a local newsgroup of some hundred users. To understand how crucial is the data size we need to understand how each algorithm deals with the memory size (space complexity) and the number of operations performed to cluster a set of data (time complexity). Table 1 shows both of these metrics for the algorithms discussed. Here, k denotes the number of clusters, t the number of iterations, and n the size of the data being clustered. It is obvious that the problem is with the $O(n^2)$ algorithms specially when n is large. Xu and Wunsch (2005) compared the time and space complexities of these algorithms and provided an additional algorithms that can handle very large datasets such as CLARA, CLARANS and BIRCH.

Table 1 Space and time complexities for clustering algorithms

<i>Algorithm name</i>	<i>Space complexity</i>	<i>Time complexity</i>
Single-link	$O(n^2)$	$O(kn^2)$
Average-link	$O(n^2)$	$O(kn^2)$
MST single-link	$O(n^2)$	$O(n^2)$
K-means	$O(n+k)$	$O(tkn)$
ROCK	$O(n^2)$	$O(n^2 \log(n))$
DBSCAN	$O(n^2)$	$O(n^2)$ or $O(n \log(n))$ with R*-tree

It is obvious that hierarchical clustering algorithms are not suitable for large datasets because of their complexities. The K-means is the most efficient algorithm among them as the complexity is almost linear (Xu and Wunsch, 2005), but it cannot handle categorical data which is very important when clustering the web. DBSCAN can be improved by using spatial indices on the data points such as R*-tree that will reduce the time complexity for it from $O(n^2)$ to $O(n \log(n))$ and generates more efficient queries (Xu and Wunsch, 2005). It is important to mention that indexing spatial data faces difficulties in high dimensions and this subject is an active area of research (Marmanis and Babenko, 2009).

6.2 High dimensionality

The world that we deal with is of three-dimensionality and if we want to cluster worlds of higher dimensionalities we need to know that these worlds are governed by different rules and different proximities (Marmanis and Babenko, 2009). Actually higher dimensionality means larger computation which will slow the algorithm down.

High dimensionality produces a problem in data separation as the distance between the point and its nearest neighbour has no difference than the distance from that point to other points when the dimensionality is high enough (Xu and Wunsch, 2008). The ‘curse of dimensionality’ is a problem that is related to high dimensionality. The term was introduced by Bellman to indicate the exponential growth of complexity in a high dimensionality situation (Xu and Wunsch, 2005), which indicates that the distance between any set of points in high dimensions are the same. In such situation there will be no effect for clustering algorithms that are based on distance measurements. Aggarwal (2002) provided a solution to this problem in.

7 Related works

Many algorithms for community detection have been proposed in the past. Newman and Girvan (2004) proposed a method for discovering communities based on hierarchical divisive algorithms where an edge is removed iteratively from the network to split it into communities. Divisive algorithms were rarely used for community detection at that time as most of the studies were based on agglomerative algorithms. The main idea of their algorithm is to remove the edge with the highest betweenness. One fast method to measure edge betweenness is the shortest-path betweenness by measuring all the shortest paths passing through a giving link. Once an edge removed a recalculation of the edge betweenness is needed for all edges which leads to high computational costs. The process continues until the final communities consist of single nodes.

Fortunato et al. (2004) implemented a hierarchical clustering algorithm based on the work of Newman and Girvan (2004). Their work was based on using centrality measure to iteratively find and remove edges in the network. Their work shows that even the algorithm runs in $O(n^4)$, it is powerful when dealing with mixed and hard detectable communities.

Newman (2004) also proposed another qualitative method for identifying communities called modularity. The *modularity* method uses a quantity function Q to measure if a specific division is meaningful when identify communities. The algorithm used was simple with reasonable running time than the other proposed algorithms at that time with worst time complexity of $O(n^2)$ on sparse network of nodes n . This method has become very popular and widely used. However, it has limits as claimed by Fortunato and Barthelemy (2007) who showed in their paper that modularity has scale that depends on the size of the network and any module that is smaller than the scale might not be determined. Li et al. (2008) proposed a quantitative modularity density measure to get over the limits of the regular modularity method proposed by Newman and Girvan. They used both the nodes and the edges in their method and showed that the optimisation of the quantitative modularity density measure will not affect the network division process which provides better detection for communities. But the method is still NP-hard.

Clauset et al. (2004) proposed another algorithm called the CNM algorithm. The CNM algorithm is a bottom-up agglomerative clustering method that uses greedy techniques to combine and merge clusters in a network. The algorithm works similarly as in Newman (2004) and gives similar results for the communities found. But it is more efficient in identifying communities since its time performance in worst case drops to $O(m * d * \log n)$.

8 Case study and evaluation

In this section, we use a case study to further explain clustering algorithms. We will explain the scripting language we have used and how we setup the environment to run the algorithms and obtain the results. Also we will discuss the results obtained using each algorithm.

The issue of identifying articles of similar topics is of great potential in the intelligent web environment, so in our case study we will use clustering algorithms to help in grouping similar articles. We collected the data from Delicious.com (2012), which is a social bookmarking service that allows users to share, store and discover web bookmarks. Since we are dealing with categorical data and keywords, represented by the articles titles, we will use ROCK and DBSCAN algorithms to define the different clusters and group the similar titles together.

8.1 Data collection

For our two experiments, we have collected a list of 48 titles for different articles from Delicious.com and saved them in CSV file. For each title, we assigned a unique ID and a username of the person who bookmarked that title. Two or more users can bookmark the same title. A sample of the dataset (11 out of 48 titles) is illustrated in Table 2.

Table 2 Sample dataset collected from Delicious.com

<i>ID</i>	<i>Name</i>	<i>Title</i>
776	user01	Google sites to add social networking in 'layers'
774	user01	40 Best And Highly Useful Websites For Adobe Photoshop Tutorials
740	user01	Nikon D7000: Camera Road Test With Chase Jarvis Chase Jarvis Blog
770	user01	Twitter is NOT a Social Network, Says Twitter Exec
722	user01	An open source collaborative network
744	user02	Google sites to add social networking in 'layers'
710	user02	40 Best And Highly Useful Websites For Adobe Photoshop Tutorials
730	user03	Google sites to add social networking in 'layers'
777	user03	40 Best And Highly Useful Websites For Adobe Photoshop Tutorials
756	user03	An open source collaborative network
733	user03	How To Discover Your Money Making Niche

8.2 Setting up the environment

The two algorithms are implemented in Java language and to execute and debug them we used BeanShell, which is a free Java interpreter. The latest Java JDK and Apache Ant should be installed in order for the BeanShell interpreter to work correctly. All of the code commands were executed through the command line in Windows OS environment.

8.3 ROCK algorithm

First, we used the ROCK algorithm to cluster the dataset. The algorithm uses the Jaccard coefficient to measure the similarities between different titles. It compares the common terms or keywords in the titles and based on that similarity titles are grouped together.

To start the experiment, we have loaded Delicious titles using the 15 most common terms only and stored them in an array. The ROCK algorithm invoked to cluster the dataset with a minimum number of clusters equal to 5. This parameter will allow the ROCK algorithm to stop before grouping all the data in one cluster. The threshold of 0.2 is used to represent the needed proximity between two points to be linked. Algorithm 6 represents the code used to execute the algorithm and print the result.

Algorithm 6 ROCK algorithm execution code

```

1 DeliciousDataset ds = DeliciousData.createDataset(15);
2 DataPoint[] dps = ds.getData();
3 ROCKAlgorithm rock = new ROCKAlgorithm(dps, 5, 0.2);
4 Dendrogram dnd = rock.cluster();
5 dnd.print(16);

```

Table 3 ROCK algorithm results

<i>Clusters for: level-16, goodness = 1.973889261532508</i>		
<i>Cluster no.</i>	<i>ID</i>	<i>Title</i>
1	799	Nikon D7000: Camera Road Test With Chase Jarvis Chase Jarvis Blog
1	688	Nikon D7000: Camera Road Test With Chase Jarvis Chase Jarvis Blog
2	708	40 Best And Highly Useful Websites For Adobe Photoshop Tutorials
2	774	40 Best And Highly Useful Websites For Adobe Photoshop Tutorials
2	710	40 Best And Highly Useful Websites For Adobe Photoshop Tutorials
3	722	An open source collaborative network
3	715	An open source collaborative network
4	520	Twitter is NOT a Social Network, Says Twitter Exec
4	566	Twitter is NOT a Social Network, Says Twitter Exec
4	744	Google sites to add social networking in 'layers'
4	730	Google sites to add social networking in 'layers'
4	776	Google sites to add social networking in 'layers'
4	770	Twitter is NOT a Social Network, Says Twitter Exec
5	740	Nikon D7000: Camera Road Test With Chase Jarvis Chase Jarvis Blog
5	720	Nikon D7000: Camera Road Test With Chase Jarvis Chase Jarvis Blog
6	777	40 Best And Highly Useful Websites For Adobe Photoshop Tutorials
6	795	40 Best And Highly Useful Websites For Adobe Photoshop Tutorials
7	681	Twitter is NOT a Social Network, Says Twitter Exec
7	500	Twitter is NOT a Social Network, Says Twitter Exec
7	790	Twitter is NOT a Social Network, Says Twitter Exec
7	780	Google sites to add social networking in 'layers'
8	735	40 Best And Highly Useful Websites For Adobe Photoshop Tutorials
8	726	40 Best And Highly Useful Websites For Adobe Photoshop Tutorials

The results of our experiment for the ROCK algorithms at level 16 shows eight clusters, (Table 3). We noticed that the algorithm clustered similar titles such as title ID 799 and title ID 688 together. On the other hand, there are articles with similar titles but the algorithm did not merge them in one cluster such as title ID 520 that is in cluster 4 and title ID 681 that is in cluster 7. The algorithm also defined the non-obvious clusters such as the titles in cluster 4; which contains different titles that are grouped together because they contain similar terms related to ‘social network’ topic. The ROCK algorithm will compare titles based on important keywords in these titles.

8.4 DBSCAN algorithm

We applied the DBSCAN algorithm to the same dataset. The algorithm used also the Jaccard coefficient to measure the similarities between different titles. To start the experiment we have loaded Delicious titles using the most 15 common terms only and stored them in an array. A Cosine distance is used as a distance metric. The DBSCAN algorithm invoked to cluster the dataset with a *distance metric*, *Eps (neighbour threshold)*, *minPoints* and *the term frequency*. Algorithm 7 represents the code used to execute the algorithm and print the result.

Algorithm 7 DBSCAN algorithm execution code

```

1  DeliciousDataset ds = DeliciousData.createDataset(15);
2  DataPoint[] dps = ds.getData();
3  CosineDistance cosD = new CosineDistance();
4  DBSCANAlgorithm dbscan = new DBSCANAlgorithm(dps,cosD,0.7,2,true);
5  dbscan.cluster();

```

The results of our experiment for the DBSCAN algorithms are shown in Table 4. The results were more accurate than the ROCK algorithm results. All similar titles are clustered together such as clusters 2 and 3 as you can notice the titles are exactly similar to each other. Non-similar titles in cluster 1 (title ID 776, title ID 566) and cluster 4 (title ID 722, title ID 711) also defined by the algorithm, where in cluster 1 it grouped the titles based on the keyword ‘social networks’ and in cluster 4 it grouped all the titles related to ‘open source’ topic. The algorithm also was able to recognise the noise elements where these points do not belong to any cluster.

8.5 Results

From both experiments, the DBSCAN and the ROCK algorithms produce good clustering results for the dataset. We noticed that the DBSCAN advanced the ROCK algorithm as it able to find the correct clusters and outliers as shown in Table 4. The ROCK is based on measuring similarity between two clusters as it finds the common neighbours for the two clusters. But relying on similarity might make the algorithm merge two clusters due to closeness even if they contain outliers or noise.

Table 4 DBSCAN algorithm results

<i>DBSCAN clustering with NeighbourThreshold = 0.7 minPoints = 2</i>		
<i>Cluster no.</i>	<i>ID</i>	<i>Title</i>
1	776	Google sites to add social networking in 'layers'
1	566	Twitter is NOT a Social Network, Says Twitter Exec
1	744	Google sites to add social networking in 'layers'
1	780	Google sites to add social networking in 'layers'
1	790	Twitter is NOT a Social Network, Says Twitter Exec
1	500	Twitter is NOT a Social Network, Says Twitter Exec
1	730	Google sites to add social networking in 'layers'
1	770	Twitter is NOT a Social Network, Says Twitter Exec
1	681	Twitter is NOT a Social Network, Says Twitter Exec
1	520	Twitter is NOT a Social Network, Says Twitter Exec
2	774	40 Best And Highly Useful Websites For Adobe Photoshop Tutorials
2	708	40 Best And Highly Useful Websites For Adobe Photoshop Tutorials
2	777	40 Best And Highly Useful Websites For Adobe Photoshop Tutorials
2	726	40 Best And Highly Useful Websites For Adobe Photoshop Tutorials
2	795	40 Best And Highly Useful Websites For Adobe Photoshop Tutorials
2	735	40 Best And Highly Useful Websites For Adobe Photoshop Tutorials
2	710	40 Best And Highly Useful Websites For Adobe Photoshop Tutorials
3	740	Nikon D7000: Camera Road Test With Chase Jarvis Chase Jarvis Blog
3	530	Nikon D7000: Camera Road Test With Chase Jarvis Chase Jarvis Blog
3	688	Nikon D7000: Camera Road Test With Chase Jarvis Chase Jarvis Blog
3	685	Nikon D7000: Camera Road Test With Chase Jarvis Chase Jarvis Blog
3	799	Nikon D7000: Camera Road Test With Chase Jarvis Chase Jarvis Blog
3	720	Nikon D7000: Camera Road Test With Chase Jarvis Chase Jarvis Blog
4	722	An open source collaborative network
4	590	An open source collaborative network
4	711	XWiki -Open Source Wiki and Content-Oriented Application Platform
4	600	An open source collaborative network
4	715	An open source collaborative network
4	756	An open source collaborative network
5	690	Apple: Sorry, Steve Jobs Isn't a Ninja
5	736	Apple: Sorry, Steve Jobs Isn't a Ninja
6	499	How To Discover Your Money Making Niche
6	733	How To Discover Your Money Making Niche
7	743	How To Create WordPress Themes From Scratch Part 1
7	533	How To Create WordPress Themes From Scratch Part 3b
7	694	How To Create WordPress Themes From Scratch Part 2
7	510	How To Create WordPress Themes From Scratch Part 3a

Table 4 DBSCAN algorithm results (continued)

<i>DBSCAN clustering with NeighbourThreshold = 0.7 minPoints = 2</i>		
<i>Cluster no.</i>	<i>ID</i>	<i>Title</i>
Noise	540	iPhone SDK 3.0 Playing with Map Kit -ObjectGraph Blog
Noise	742	This Is The Second Time A Google Engineer Has Been Fired For Accessing User Data
Noise	577	Enhance your web forms with new HTML5 features
Noise	745	Resize or Move a Complete Flash Animation in One Go
Noise	746	10 Things You Didn't Know About the New #Twitter/via @gigaom #news #sm
Noise	732	How To Handle Customers During Virtual Assistant Problems
Noise	705	article on social media ad campaigns
Noise	587	The Business Plan
Noise	791	CSS Color Names
Noise	601	Typography: Web Style Guide 3
Noise	753	in 4 US Adults Now Use Mobile Apps [STATS]

9 Conclusions

Clustering algorithms are an important approach to divide and analyse data. There are many different types of clustering each with its own technique. We can apply them on many phenomena in this world actually any dataset that consists of elements are qualified for applying clusters. We have discussed six clustering algorithms in this paper: single link, average link, MST single link, K-means, ROCK, and DBSCAN. The discussion was based on how accurate is to apply them on the social web.

The single-link, average-link, and MST single-link algorithms are agglomerative hierarchical algorithms. They do not perform efficiently (both with respect to time and space) on large datasets even though they are easily implemented algorithms.

The K-means algorithm is a partitional algorithm, which is more efficient than link-based algorithms. But it does not work with categorical data since it relies on the idea of centroid. Moreover it cannot handle outliers (the points that are far away from the main clusters) (Xu and Wunsch, 2005).

The ROCK algorithm is a hierarchical agglomerative algorithm that can handle categorical data since it relies on the links as measures more than the distance. But it has high time and space complexities.

The DBSCAN algorithm is a density-based algorithm that uses point density to identify clusters in a space. It can handle outliers even though its time and space complexity are high.

The algorithms discussed in this paper are used for the identification of groups of users and data on a website. We can combine algorithms such as the K-means with other algorithms. K-means is the preferred to use since it is simple and fast and can run on parallel computational platforms. Combining different algorithms together will maximise the benefits of these algorithms and guarantees the quality of resulted clusters (Kotsiantis and Pintelas, 2004). One example would be combining the efficient K-means algorithm

with the powerful ROCK algorithm (if the data is Boolean or categorical) or DBSCAN algorithm (if the data is spatial). One scenario would be using K-means on the high-level clusters then process them with the ROCK or the DBSCAN algorithm.

References

- Adomavicius, G. and Tuzhilin, A. (2005) 'Toward the next generation of recommender systems: a survey of the state-of-the-art and possible extensions', *IEEE Transactions on Data and Knowledge Engineering*, Vol. 17, No. 6, pp.734–749.
- Aggarwal, C.C. (2002) 'Towards meaningful high-dimensional nearest neighbour search by human-computer interaction', *Proceedings of the 18th International Conference on Data Engineering*.
- Arthur, D. and Vassilvitskii, S. (2007) 'K-means++: the advantages of careful seeding', in *SODA '07: Proceedings of the Eighteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, pp.1027–1035.
- Berkhin, P. (2002) 'Survey of clustering data mining techniques', Technical report, Accrue Software, San Jose, CA.
- Clauset, A., Newman, M.E.J. and Moore, C. (2004) 'Finding community structure in very large networks', *Physical Review E*, Vol. 70, No. 6, 066111.
- DBSCAN (2012) *Wikipedia* [online] <http://en.wikipedia.org/wiki/DBSCAN> (accessed 2 April 2012).
- Delicious.com (2012) *Wikipedia* [online] <http://en.wikipedia.org/wiki/Delicious.com> (accessed 2 April 2012).
- Ester, M., Kriegel, H-p., Jörg, S. and Xu, X. (1996) 'A density-based algorithm for discovering clusters in large spatial databases with noise', *Proceedings of 2nd International Conference on Knowledge Discovery and Data Mining (KDD-96)*.
- Everitt, B.S., Landau, S. and Leese, M. (2001) *Cluster Analysis*, 4th ed., Wiley, West Sussex, UK.
- Discovering Who To Follow (2010) *Twitter Blog* [online] <http://blog.twitter.com/2010/07/discovering-who-to-follow.html> (accessed 2 April 2012).
- Fortunato, S. and Barthelemy, M. (2007) 'Resolution limit in community detection', *Proc. Natl. Acad. Sci. USA*, Vol. 104, No. 1, pp.36–41.
- Fortunato, S., Latora, V. and Marchiori, M. (2004) 'Method to find community structures based on information centrality', *Physical Review E*, Vol. 70, No. 5, 056104.
- Gower, J.C. (1971) 'A general coefficient of similarity and some of its properties', *Biometrics*, Vol. 27, No. 4, pp.857–871.
- Hammouda, K.M. (2002) 'Web mining: identifying document structure for web document clustering', Master's thesis, University of Waterloo, Waterloo, Ontario, Canada.
- Haveliwala, T.H., Gionis, A. and Indyk, P. (2000) 'Scalable techniques for clustering the web', *ACM Computing Surveys*, Vol. 31, No. 3, pp.264–323.
- Hruschka, E.R., Campello, R.J.G.B., Freitas, A.A. and Leon, P. (2009) 'A survey of evolutionary algorithms for clustering', *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)*, Vol. 39, No. 2, pp.133–155.
- Jain, A.K. and Dubes, R.C. (1988) *Algorithms for Clustering Data*, Prentice-Hall, Inc., Upper Saddle River, NJ, USA.
- Jain, A.K., Murty, M.N. and Flynn, P.J. (1999) 'Data clustering: a review', *ACM Computing Surveys*, Vol. 31, No. 3, pp.264–323.
- Kanungo, T., Mount, D.M., Netanyahu, N.S., Piatko, C., Silverman, R. and Wu, A.Y. (2002) 'An efficient k-means clustering algorithm: analysis and implementation', *IEEE Transactions on Pattern Analysis and Machine Intelligence*, Vol. 24, No. 7, pp.881–892.

- Kotsiantis, S.B. and Pintelas, P.E. (2004) 'Recent advances in clustering: a brief survey', *WSEAS Transactions on Information Science and Applications*, Vol. 1, No. 1, pp.73–81.
- Li, Z., Zhang, S., Wang, R-S., Zhang, X-S. and Chen, L. (2008) 'Quantitative function for community detection', *Phys. Rev. E*, Vol. 77, No. 3, 036109.
- Marmanis, H. and Babenko, D. (2009) *Algorithms of the Intelligent Web*, Manning Publications Co, Greenwich, Connecticut, USA.
- Mishra, N., Schreiber, R., Stanton, I. and Tarjan, R. (2007) 'Clustering social networks', in Bonato, A. and Chung, F.R.K. (Eds.): *Algorithms and Models for the Web-Graph, Lecture Notes in Computer Science*, Chapter 5, Vol. 4863, pp.56–67, Springer, Berlin, Heidelberg.
- Mocian, H. (2009) 'Survey of distributed clustering techniques', MSc Internal Research Project, Computer Science, Imperial College London.
- Naughton, M., Kushmerick, N. and Carthy, J. (2006) 'Clustering sentences for discovering events in news articles', *Advances in Information Retrieval*, Vol. 3936, pp.535–538.
- Newman, M.E.J. (2004) 'Fast algorithm for detecting community structure in networks', *Phys. Rev. E*, Vol. 69, No. 6, 066133.
- Newman, M.E.J. and Girvan, M. (2004) 'Finding and evaluating community structure in networks', *Phys. Rev. E*, Vol. 69, No. 2, 026113.
- Rajeev, S.G., Rastogi, R. and Shim, K. (1999) 'ROCK: a robust clustering algorithm for categorical attributes', *15th International Conference on Data Engineering*.
- SemanticWeb (2010) *Wikipedia* [online] http://en.wikipedia.org/wiki/Semantic_Web (accessed 2 April 2012).
- Stein, B. and Busch, M. (2005) 'Density-based cluster algorithms in low dimensional and high-dimensional application', *Second International Workshop on Text-Based Information Retrieval*.
- Tan, P-N., Steinbach, M. and Kumar, V. (2005) *Introduction to Data Mining*, 1st ed., Addison Wesley, Boston, MA, USA.
- Wu, B.Y. and Chao, K-M. (2004) *Spanning Trees and Optimization Problems*, Chapman and Hall/CRC Press, USA.
- Xu, R. and Wunsch, D. (2008) *Clustering*, Wiley-(IEEE Press Series on Computational Intelligence), Hoboken, New Jersey, USA.
- Xu, R. and Wunsch, I. (2005) 'Survey of clustering algorithms', *IEEE Transactions on Neural Networks*, Vol. 16, No. 3, pp.645–678.