Identifying Single Clusters in Large Data Sets

Frank Klawonn
University of Applied Sciences Braunschweig/Wolfenbuettel, Germany

Olga Georgieva
Institute of Control and System Research, Bulgaria

INTRODUCTION

Most clustering methods have to face the problem of characterizing good clusters among noise data. The arbitrary noise points that just do not belong to any class being searched for are of a real concern. The outliers or noise data points are data that severely deviate from the pattern set by the majority of the data, and rounding and grouping errors result from the inherent inaccuracy in the collection and recording of data. In fact, a single outlier can completely spoil the least squares (LS) estimate and thus the results of most LS based clustering techniques such as the hard C-means (HCM) and the fuzzy C-means algorithm (FCM) (Bezdek, 1999).

For these reasons, a family of robust clustering techniques has emerged. There are two major families of robust clustering methods. The first includes techniques that are directly based on robust statistics. The second family, assuming a known number of clusters, is based on modifying the objective function of FCM in order to make the parameter estimates more resistant to the data noise. Among them one promising approach is the noise clustering (NC) technique (Dave, 1991; Klawonn, 2004). It maintains the principle of probabilistic clustering, but an additional noise cluster is introduced. NC was developed and investigated in the context of a variety of objective function-based clustering algorithms and it has demonstrated its reliable ability to detect clusters amongst noise data.

BACKGROUND

Objective function-based clustering aims at minimizing an objective function that indicates a kind of fitting error of the determined clusters to the given data. In this objective function, the number of clusters has to be fixed in advance. However, as the number of clusters is usually unknown, an additional scheme has to be applied to determine the number of clusters (Guo, 2002; Tao, 2002). The parameters to be optimized are the membership degrees that are values of belonging of each data point to every cluster, and the parameters, characterizing the cluster, which finally determine the distance values. In the simplest case, a single vector named cluster centre (prototype) represents each cluster. The distance of a data point to a cluster is simply the Euclidean distance between the cluster centre and the corresponding data point. More generally, one can use the squared inner-product distance norm, in which by a norm inducing symmetric and positive matrix different variances in the directions of the coordinate axes of the data space are accounted for. If the norm inducing matrix is the identity matrix we obtain the standard Euclidean distance that form spherical clusters. Clustering approaches that use more complex cluster prototypes than only the cluster centres, leading to adaptive distance measures, are for instance the Gustafson-Kessel (GK) algorithm (Gustafson, 1979), the volume adaptation strategy (Höppner, 1999; Keller, 2003) and the Gath-Geva (GG) algorithm (Gath, 1989). The latter one is not a proper objective function algorithm, but corresponds to a fuzzified expectation maximization strategy. No matter which kind cluster prototype is used, the assignment of the data to the clusters is based on the corresponding distance measure. In hard clustering, a data object is assigned to the closest cluster, whereas in fuzzy clustering a membership degree, that is, a value that belongs to the interval [0,1] is computed. The highest membership degree of a data corresponds to the closest cluster.

Noise clustering has a benefit of the collection of the noise points in one single cluster. A virtual noise prototype with no parameters to be adjusted is introduced that has always the same distance to all points in the data set. The remaining clusters are assumed to be the good clusters in the data set. The objective function that considers the noise cluster is defined in the same manner as the general scheme for the clustering minimization functional. The main problem of NC is the proper choice of the noise distance. If it is set too small, then most of the points will get classified as noise points, while for a large noise distance most of the points will be classified into clusters other than the noise cluster. A right selection of the distance will result in a classification where the points that are actu-

ally close enough to the good clusters will get correctly classified into a good cluster, while the points that are away from the good clusters will get classified into the noise cluster. The detection of noise and outliers is a serious problem and has been addressed in various approaches (Leski, 2003; Yang, 2004; Zhang, 2003). However, all these methods need complex additional computations.

MAIN THRUST

The purpose of most clustering algorithms is to partition a given data set into clusters. However, in data mining tasks partitioning is not always the main goal. Finding interesting patterns or substructures in a large data set in the form of one or a few clusters that do not cover or partition the data set completely is an important issue in data mining. For this purpose a new clustering algorithm named noise clustering with one good cluster**,** based on the noise clustering technique and able to detect single clusters step-by-step in a given data set, has been recently developed (Georgieva, 2004). In addition to identifying clusters step by step, as a side-effect noise data are detected automatically.

The algorithm assesses the dynamics of the number of points that are assigned to only one good cluster of the data set by slightly decreasing the noise distance. Starting with some large enough noise distance it is decreased with a prescribed decrement till the reasonable smallest distance is reached. The number of data belonging to the good cluster is calculated for every noise distance using the formula for the hard membership values or fuzzy membership values, respectively. Note that in this scheme only one cluster centre has to be computed, which in the case of hard noise clustering is the mean value of the good cluster data points and in the case of fuzzy noise clustering is the weighted average of all data points. It is obvious that by decreasing the noise distance a process of "loosing" data, that is, separating them to the noise cluster, will begin. Continuing to decrease the noise distance, we will start to separate points from the good cluster and add them to the noise cluster. A further reduction of the noise distance will lead to a decreasing amount of data in the good cluster until the cluster will be entirely empty, as all data will be assigned to the noise cluster. The described dynamics can be illustrated in a curve viewing the number of data points assigned to the good cluster over the noise distance. In this curve a plateau will indicate that we are in a phase of assigning proper noise data to the noise cluster, whereas a strong slope means that we actually loose data belonging to the good cluster to the noise cluster.

Generally, a number of clusters with different shapes and densities exist in large data sets and thus a complicated dynamics of the data assigned to the single cluster will be observed. However, the smooth part of the considered curve corresponds to the situation that a relatively small amount of data is removed, which is usually caused by loosing noise data. When we loose data from a good cluster (with higher density than the noise data), a small decrease of the noise distance will lead to a large amount of data we loose to the noise cluster, so that we will see a strong slope in our curve instead of a plateau. Thus, a strong slope indicates that at least one cluster is just removed and separated from the (single) good cluster we try to find. By this way the algorithm determines the number of clusters and detects noise data. It does not depend on the initialisation, so that the danger of converging into local optima is further reduced compared to standard fuzzy clustering (Höppner, 2003).

The described procedure is implemented in a cluster identification algorithm that assesses the dynamics of the quantity of the points assigned to the good cluster or equivalently assigned to the noise cluster through the slight decrease of the noise distance. By detecting the strong slopes the algorithm separates one cluster at every algorithm pass. A significant reduction of the noise is achieved even in the first algorithm pass. The clustering procedure is repeated by proceeding with a smaller data set as the original one is reduced by the identified noise data and data belonging to the already identified cluster(s).

The curve that is determined by the dynamics of the data assigned to the noise cluster is smoother in case of fuzzy noise clustering compared to the hard clustering case due to the fuzzily defined membership values. Also local minima of this curve could be observed due to the given freedom of the points to belong to both good and noise cluster simultaneously. Fuzzy clustering can deal better with complex data sets than hard clustering due to the given relative degree of membership of a point to the good cluster. However, for the same reason the amount of the identified noise points is less than in the hard clustering case.

FUTURE TRENDS

Whereas the standard clustering partitions the whole data set, the main goal of the noise clustering with one good cluster is to identify single clusters even in the case when a large part of the data does not have any kind of group structure at all. This will have a large benefit in some application areas of cluster analysis like, for

2 more pages are available in the full version of this document, which may be purchased using the "Add to Cart" button on the publisher's webpage: www.igi-global.com/chapter/identifying-single-clusters-large-data/10664

Related Content

Classification and Regression Trees
Johannes Gehrke (2005). *Encyclopedia of Data Warehousing and Mining (pp. 141-143).*
www.irma-international.org/chapter/classification-regression-trees/10581

Rule Generation Methods Based on Logic Synthesis
Marco Muselli (2005). *Encyclopedia of Data Warehousing and Mining (pp. 978-983).*
www.irma-international.org/chapter/rule-generation-methods-based-logic/10738

Unsupervised Mining of Genes Classifying Leukemia
Diego Liberati, Sergio Bittantiand Simone Garatti (2005). *Encyclopedia of Data Warehousing and Mining (pp. 1155-1159).*
www.irma-international.org/chapter/unsupervised-mining-genes-classifying-leukemia/10771

Intelligent Cache Management for Mobile Data Warehouse Systems
Shi-Ming Huang, Binshan Linand Qun-Shi Deng (2008). *Data Warehousing and Mining: Concepts, Methodologies, Tools, and Applications (pp. 1539-1556).*
www.irma-international.org/chapter/intelligent-cache-management-mobile-data/7714

Entity Resolution on Names
(2014). *Innovative Techniques and Applications of Entity Resolution (pp. 41-66).*
www.irma-international.org/chapter/entity-resolution-on-names/103243