ANALYSIS OF SETTLEMENT STRUCTURES BY GRAPH-BASED CLUSTERING

Karl-Heinrich Anders, Monika Sester and Dieter Fritsch
Institute for Photogrammetry
Stuttgart University
Geschwister-Scholl-Strasse 24, 70174 Stuttgart, Germany
Ph.: +49-711-121-4090, Fax: +49-711-3297
e-mail: {kh.anders, monika.sester, dieter.fritsch}@ifp.uni-stuttgart.de
Semantische Modellierung; *SMATI 99*

KEY WORDS: GIS, Spatial Data Interpration, Clustering, Data Aggregation, Digital Cartography

ABSTRACT

The automatic analysis of spatial data sets presumes to have techniques for interpretation and structure recognition. Such procedures are especially needed in GIS and digital cartography in order to automate the time-consuming data update and to generate multi-scale representations of the data. In order to infer higher level information from a more detailed data set, coherent, homogeneous structures in a data set have to be delineated. There are different approaches to tackle this problem, e.g. model based interpretation, rule based aggregation or clustering procedures. In this paper, an approach for the analysis of settlement structures based on graph clustering techniques is presented.

1 INTRODUCTION AND MOTIVATION

The ever increasing amount of data and information available demands for an automation of its use. Users need adequate search tools in order to quickly access and filter relevant information. Data Mining has evolved as a branch of computer science, which tries to structure data and find inherent, possibly important, relations in the data. In general, it deals with finding facts by inference; finding information in unstructured data, or in data which is not structured explicitly for the required purpose.

Besides mere textual information, more and more the importance of image, graphical or geometric data is emerging: thus not only the attributes of an object are important, but also its spatial location. Imagine that the place of living of a person can also be used to infer his/her social status. Or the position of a good in a supermarket can influence the way it is bought. The general applications of spatial data mining are abundant: Interpretation and analysis of spatial phenomena, interpretation of images, etc. The basic tools of (Spatial) Data Mining are machine learning techniques, cluster analysis and interpretation procedures.

In GIS and digital cartography, respectively, there is a growing demand for such techniques: huge spatial data sets are being acquired and have to be kept up to date at ever increasing cycles; furthermore, information of different levels of detail is required in order to compensate for the requirements of different applications. One important application is the scale dependent data representation for quick visualization on a computer screen. In cartography, typically the data of different scales are acquired, managed and updated separately – a highly time consuming and labor intensive task. In order to accelerate update cycles and deliver actual information on-the-fly, tools and techniques for automation of initial data capture and update are required.

In Germany, the digital cadastral map (ALK) in scale 1:1,000 and the digital topographic database in scale 1:25,000 (ATKIS 25) are acquired totally separately. Therefore, also the updates have to be carried out in both data sets separately. If automatic generalization tools were available, the less detailed data set could be derived from the higher detailed one. This is envisioned by the german surveying agency in a project to derive the scale 1:50,000. There is a decision to derive this scale totally automatically from the scale 1:25,000. In order to do so, generalization or aggregation procedures and interpretation tools are needed. In this context, the automatic generation of settlement structures is a big challenge. Starting from a highly detailed representation including individual buildings and their respective attributes, an interpretation strategy has to automatically derive settlement areas. This task corresponds to the problem of finding a higher level structure in a collection of lower level details.

Figure 1 shows an example of a data set used in this study and the desired result: the regions which are outlined manually show similar characteristics both in geometric structure (local density) and in semantics

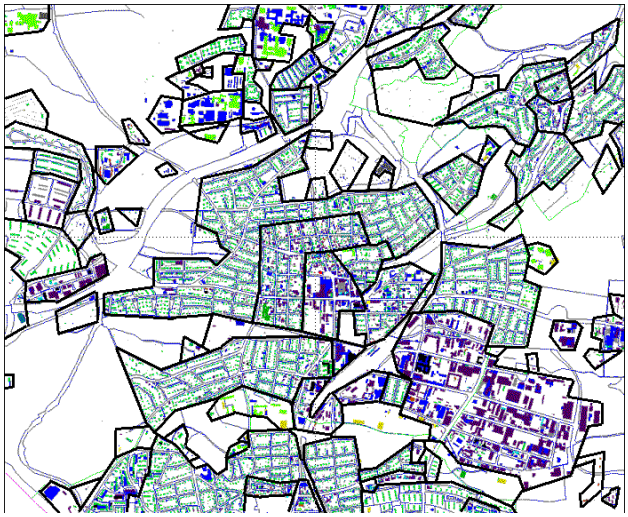

Figure 1: Building ground plans and settlement structures

(here indicated by different colors / shades of gray). The task is to automate this process and find meaningful clusters automatically.

An earlier approach [Anders & Sester 1997] focused on the modeling of the spatial situation in a semantic network and an explicit provision of a set of rules of how to aggregate objects. The prerequisite is the availability of an explicit and complete model of the situation and of the aggregation rules. Such rules are often hard to find and usually also subjective. The aim of this paper is to consider the problem as a general task of finding higher level structures in a seemingly arbitrary collection of (labeled) objects. This can be transferred to the abstract problem of considering the objects and their behaviour as a stochastic point process. In this point-collection, meaningful structures have to be identified, namely homogeneous clusters. Thus the approach relies on physiological observations of humans: humans use spatial neighborhood relations in order to find gestalt objects and separate objects from background.

Homogeneity here is considered both concerning geometry, i.e. point density, and concerning semantics, i.e. thematic 'density', namely similarity. Ideally, data mining approaches do not rely on any prior information, e.g. thresholds or parameters, which tune the process. In cluster analysis, usually the number of clusters or an information about the statistical distribution of the data is required. This approach focuses on procedures which are most generally applicable (independent on the type of objects) and need no or only few parameters. Furthermore it is important that arbitrary cluster forms can be identified, when no prior knowledge about the objects is assumed to be known. Such tasks can be tackled by clustering processes – the important prerequisite is the modeling of the neighborhood, which can be achieved by neighborhood graphs.

Thus, in this paper the notion of neighborhood is extended: in the original research project, neighborhood was defined by topologic neighborhood, i.e. objects sharing a common edge. Here neighborhood is defined object dependent in the clustering process. Our procedure requires the following steps: transformation of the data into a point-process, clustering of the points based on neighborhood graphs, delineation of the cluster shapes and finally derivation of cluster characteristics. These characteristic features can be used for object identification and classification.

The paper is structured as follows: After a short review of related work, the theoretical background of our concept, namely the clustering based on relational graphs is presented: The individual steps for clustering of point processes are identified, i.e. the establishment of the neighborhood graph, the clustering based on geometric (density) and thematic aspects, as well as the delineation of the cluster-shapes. After the presentation of the conceptual background, the concrete application of derivation of settlement structures is presented. First results are given and the next steps of our work are identified. Finally, an outlook on possible extensions and other applications concludes the paper.

2 RELATED WORK

In the context of data aggregation, there are many approaches in GIS and in digital cartography, namely in model or database generalization. [Richardson 1996] and [van Smaalen 1996] present approaches to come from one detailed scale to the next based on a set of rules. If such rules are known or models of the situation are available, good results can be achieved (cf. [Sester, Anders & Walter 1998]). However, the main problem being the definition of the rules and the control strategy to infer new data from it [Ruas & Lagrange 1995]. Current concepts try to integrate learning techniques for the derivation of the necessary knowledge [Plazanet, Bigolin & Ruas 1998], [Sester 1999].

Clustering is a well established technique for data interpretation. It usually requires prior information, e.g. about the statistical distribution of the data or the number of clusters to detect. Existing clustering algorithms, such as k-means [Jain & Dubes 1988], PAM [Kaufman & Rousseeuw 1990], CLARANS [Ng & Han 1994], DBSCAN [Ester, Kriegel, Sander & Xu 1996], CURE [Guha, Rastogi & Shim 1998], and ROCK [Guha, Rastogi & Shim 1999] are designed to find clusters that fit some static models. For example, k-means, PAM, and CLARANS assume that clusters are hyper-ellipsoidal or hyper-spherical and are of similar sizes. The DBSCAN algorithm assumes

that all points of a cluster are *density reachable* [Ester et al. 1996] and points belonging to different clusters are not. All these algorithms can breakdown if the choice of parameters in the static model is incorrect with regarding to the data set being clustered, or the model did not capture the characteristics of the clusters (e.g. shapes, sizes, densities). In the following, we give a brief overview of existing clustering algorithms.

2.1 Non-hierarchical Schemes

Non-hierarchical clustering techniques are also called partitional clustering techniques. These approaches attempt to construct a simple partitioning of a data set into a set of k non-overlapping clusters such that the partitions optimize a given criterion. Each cluster must contain at least one data element, and each data element must belong to exactly one group. In most of the partitional methods an initial partitioning is chosen and then the cluster membership is changed in order to obtain a better partitioning. *Centroid based* methods like the k-means method [MacQueen 1967], [Jain & Dubes 1988] and the ISODATA [Ball & Hall 1965] method try to assign data elements to clusters such that the mean square distance of data elements to the centroid of the assigned cluster is minimized. These techniques are suitable only for data in metric spaces, because they have to compute a centroid of a given set of data elements. *Medoid based* approaches as CLARANS [Ng & Han 1994] and PAM [Kaufman & Rousseeuw 1990] try to find a so called medoid which is a representative data element that minimize the sum of the distances between the medoid and the data elements assigned to this medoid.

One disadvantage of centroid and medoid based methods is that not all values of k lead to natural cluster so it is useful to run the algorithm several times with different values for k to select the best partition. With a given optimization criterion this decision can be automated. The main drawback of both methods is that they will fail for data sets in which data elements belonging to a cluster are closer to the representative of another cluster than to the representative of their own cluster. This case is typical for many natural clusters if the cluster shapes are concave or their sizes vary largely.

2.2 Hierarchical Schemes

Hierarchical cluster schemes constructs a dendrogram is a tree structure which represents a sequence of nested clusters. This sequence represents multiple levels of partitioning. On the top is a single cluster which includes all other clusters. At the bottom are the data elements representing single element clusters. Dendrograms can be constructed top-down or bottom-up. The btotom-up method is known as the agglomerative approach, where each data element starts out as a seperate cluster. In each step of an agglomerative algorithm the two most similar clusters are grouped together based on similarity measures in subsequent steps and the total number of clusters is decreased by one. These steps can be repeated until one large cluster remain or a given number of clusters is optained or the distance between two closest clusters is above a certain threshold. The top-down method known as the divisive approach works in the reverse direction. Agglomerative methods seems to be the most popular in the literature.

In the literature one can find many different variations of hierarchical algorithms. Basically, these algorithms can be distinguished by their definition of similarity and how they update the similarity between existing clusters and the merged clusters. In general, the aproaches described are alternative formulations or minor variations of the following three concepts:

- centroid or medoid based methods,
- linkage based methods,
- variance or error sum of squares error.

The centroid or medoid based approaches also fail on clusters of arbitrary shapes and different sizes like non-hierarchical methods, such as k-means and k-medoid. The oldest linkage based method is the *single linkage* algorithm, sometimes referred to as the nearest neighbor approach. In the single linkage method, no representative exists. The cluster is represented by all data elements in the cluster and the similarity between two clusters is the distance between the closest pair of data elements belonging to different clusters. The single linkage method is able to find clusters of arbitrary shape and different sizes, but it will fail at poorly seperated clusters and is susceptible to noise and outliers.

In order to avoid these drawbacks algorithms like the *shared near neighbors method* [Jarvis & Patrick 1973], *CURE* [Guha et al. 1998] or *ROCK* [Guha et al. 1999] were proposed. Instead of using a single centroid to represent a cluster, CURE choose a constant number of representative points to describe a cluster. The ROCK algorithm operates on a derived similarity graph and scales the aggregate *inter-connectivity* with respect to a predefined inter-connectivity model. The shared near neighbors method use a k-nearest-neighbour graph to determine the similarity between two clusters. The advantage of this clustering method over most other alternatives is that it is independent of absolute scale.

A major limitation of existing agglomerative hierarchical schemes such as the *Group Averaging Method*

[Jain & Dubes 1988], CURE, and ROCK is that the merging decisions are based on static modeling of the clusters to be merged. More information about the limitations of existing hierarchical methods can be found in [Karypis, Han & Kumar 1999].

3 GRAPH-BASED CLUSTERING

The most powerful methods of clustering in difficult problems, which give results having the best agreement with human performance, are the graph-based methods [Jaromczyk & Toussaint 1992]. The idea is extremely simple: Compute a neighborhood graph (such as the minimal spanning tree) of the original points, then delete any edge in the graph that is much longer (according to some criterion) than its neighbors. The result is a forest and each tree in the forest represents a cluster.

3.1 Similarity

In general, hierachical cluster algorithms work implicitly or explicitly on a similarity matrix such that every element of the matrix represents the similarity between two elements. In each step of the algorithm the similarity matrix is updated to reflect the revised similarities. Basically, all these algorithms can be distinguished based on their definition of similarity and how they update the similarity matrix. In spatial clustering algorithms one can discriminate between *spatial similarity* and *semantic similarity* which means the similarity of non-spatial attributes.

Spatial Similarity implies the definition of a neighborhood concept which can be defined on geometric attributes, such as coordinate, distance, density, and shape. The computation of a spatial similarity matrix can be seen as the construction of a weighted graph, so called *neighborhood graph*, where each element is represented by a node and each neighborhood relationship (similarity) is an edge. There are efficient algorithms to compute neighborhood graphs [Jaromczyk & Toussaint 1992] which can be used to compute a spatial similarity matrix.

Whereas geometric similarity can be evaluated based on a given metric scale, in order to evaluate semantic similarity, adequate measures, so called *nominal scales* have to be established. The simplest possibility is to use identity; more sophisticated measures are the Hamming distance, Dice coefficient, Tanimoto coefficient or the Levenshtein distance. These measures, however, only work by comparing the differences of the strings. Another way is to establish compatibility matrices, which indicate the semantic closeness of objects: in the context of buildings, a garage and a residential building are more close than a garage and an industrial building. Such values can be derived empirically by evaluating the neighborhood relationships of objects in existing data sets.

In order to integrate both semantic and geometric homogeneity constraints in the clustering process, different strategies are possible:

- two-step procedure:
  1. clustering based on geometry alone using RNG (no parameters necessary)
  2. semantic clustering inside these clusters
- one-step procedure: definition of an integrated measure of spatial and semantic distance (problem to evaluate thematic against spatial distance). One simple example for an integrated measure is the mean value of the spatial and semantic similarity.

3.2 Neighborhood Graphs

A general introduction to the subject of proximity graphs is given in [Jaromczyk & Toussaint 1992]. Neighborhood graphs also called *proximity graphs* [Toussaint 1991], are used as tools in disciplines where shape and structure of point sets are of primary interest. These include for example visual perception, computer vision and pattern recognition, cartography and geography, and biology.

Relative neighborhood graphs (RNG's), introduced by [Toussaint 1980], capture proximity between points by connecting nearby points with a graph edge. The many possible notions of *nearby* (in several metrics) lead to a variety of related graphs. It is easiest to view the graphs as connecting points only when certain regions of space are empty. Some typical neighborhood graphs are

- Relative neighborhood graph [Toussaint 1980],
- Gabriel graph [Gabriel & Sokal 1969],
- β-skeleton [Kirkpatrick & Radke 1985],
- Sphere of influence graph [Toussaint 1988],
- α-graphs [Edelsbrunner, Kirkpatrick & Seidel 1983].

3.3 Shape and Characteristics of Clusters

After the cluster generation, the characteristics of the clusters can be calculated. These characteristics are geometric features, like size, shape, average density, etc. In addition to these unary features, also binary features – relations – between the clusters can be used. These characteristics can then help either to identify clusters with similar characteristics, or even to

Figure 2: Starting situation: buildings given in polygonal representation

identify objects. This can be achieved by well known pattern recognition or interpretation procedures, e.g. model based interpretation. In this way, e.g. dense inner city centers can be differentiated from rural or industrial areas.

The computation of the *shape* of a point set is a non-trivial problem. In contrast to other geometric notions, such as diameter, volume, or convex hull the geometric notion of shape has no associated formal meaning [Edelsbrunner & Mücke 1994]. A fair amount of related work has been done for planar point sets, and some for three dimensional point sets. One of the first who considered the problem of computing the shape of a point set as a generalization of the convex hull was [Jarvis 1977]. A general and mathematically well defined concept of shape, called α-*shapes* was introduced by [Edelsbrunner et al. 1983] and was generalized to three dimensions in [Edelsbrunner & Mücke 1994]. The α-shapes has been used for cluster analysis, molecular modeling, and the analysis of medical data, among other applications.

4 CLUSTERING OF SETTLEMENT STRUCTURES

In order to separate meaningful settlement structures, different steps have to be taken. Starting point is a typical spatial data set of buildings, given in terms of polygonal object description (cf. Figure 2). The task is to identify coherent, homogeneous regions in the data set.

4.1 Preprocessing

In order to apply the above mentioned clustering techniques, the data has to be transformed into a initial similarity matrix, which will be done by computing the relative neighborhood graph of the given building groundplans. This graph structure consists of nodes which represent the building groundplans and edges representing the spatial relations between the building groundplans. Given a similarity matrix, many methods can be used to find a graph representation [Jarvis & Patrick 1973], [Jain & Dubes 1988], [Guha et al. 1999], [Karypis et al. 1999]. Modeling data items as a graph is very common in many hierarchical clustering algorithms. Algorithms based on the single linkage, complete linkage, or group average method [Jain & Dubes 1988] operate on a complete graph. Other methods like the shared near neighbours approach or CHAMELEON are using the k-nearest-neighbor graph which is a sparse graph.

In our approach we use the relative neighborhood graph 5. There are three reasons for using the RNG.

- The RNG is a sparse graph.
- The RNG can be computed using the Delaunay triangulation (DT).
- The RNG defines the notion of neighborhood dynamically. Independent from a global neighborhood density value this method computes a natural form of neighborhood. For example the DBSCAN [Ester et al. 1996] algorithm needs to specify a global neighbor density.

Starting with the groundplans, the DT can be calculated in several ways which basically lead to comparable results. A straightforward method is to approximate the objects by their centroid and apply a Delaunay triangulation on this collection of points. The other option is to introduce each object point into the triangulation, and to use a constrained triangulation, in order to force the object edges to be edges of the resulting triangulation. This structure has to be analyzed with respect to the relations of the objects to each other: an object shares relations to all other objects it is linked to by a common edge. The weight of the edge (the relation) can be derived from all the given edges two objects share either as minimum, maximum, or mean distance value. The result is a general attributed graph, where the nodes are coordinate free and, the edges represent the neighborhood, including the distances. Whereas in the first option the edge weight is derived using the distance between the nodes (centroids), it is directly given in the second case. Figures 3 and 4 show the results of the two operations. The computation of the DT has one great benefit, because the DT can be used to compute the RNG in $O(nlogn)$ time. In the euclidean metric the DT is a supergraph of the RNG [Toussaint 1980]. An algorithm for the RNG in the euclidean metric using the DT was developed by [Supowit 1983].

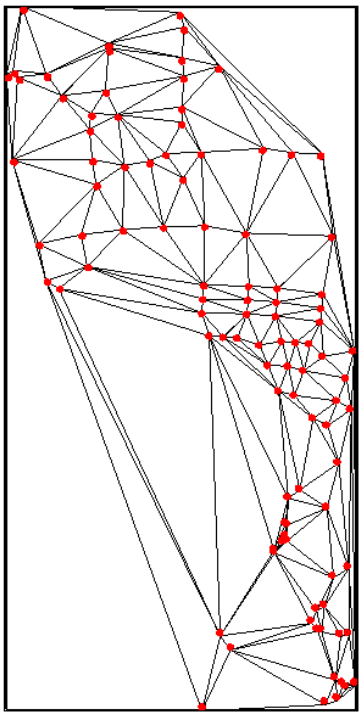

Figure 3: Triangulation of object centroids.

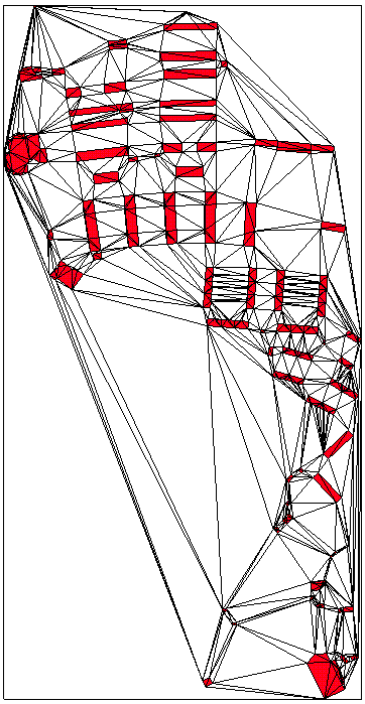

Figure 4: Constrained triangulation of complete objects.

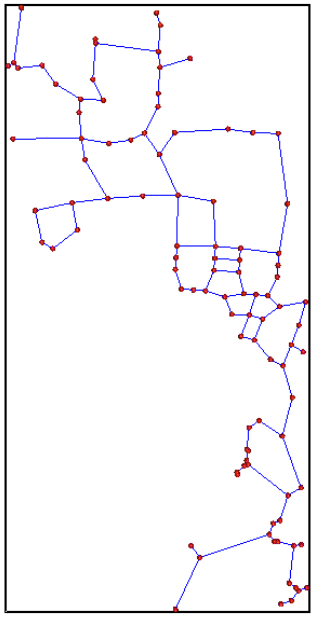

Figure 5: Relative Neighborhood Graph (RNG)

4.2 Clustering

When the preprocessing is finished we have a relative neighborhood graph 5 which is a connected graph. The following clustering process splits this connected graph in sub-graphs which represents the final clusters. This splitting has to be done in a adequate way.

In this approach we use only the spatial similarity, because the semantic similarity is not implemented yet. The splitting of the graph is based on removing outliers edges. For every node in the graph the local neighbor density is estimated. The size of the Voronoi cell around a node is the measure of the local density. Another measure which is used here is based on the Delaunay triangulation of the preprocessing step. The mean distance of all adjacent nodes is chosen as an estimator for the local density.

Methods like CHAMELEON [Karypis et al. 1999] use a two-phase clustering algorithm. First, the sparse graph is partitioned into sub-graphs of a given size and then an agglomerative hierarchical clustering algorithm is used to merge this initial sub-graphs together.

Figure 7 and 8 give another example of a larger scene.

The derived clusters are the starting point of the second step, the semantic clustering, which is however subject to our further work. This step presumes to have the thematic data available, as well as a measure for semantic similarity, which we will derive from existing data sets. Finally the cluster shapes will be outlined and their characteristics computed.

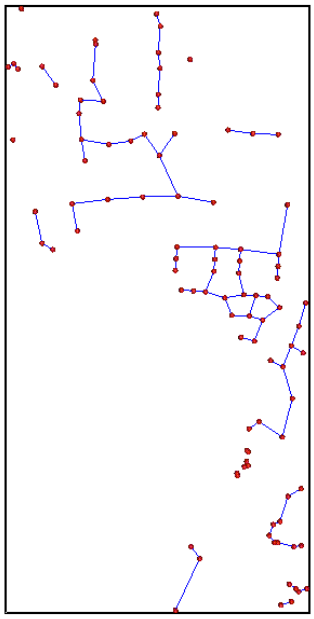

Figure 6: Result of clustering

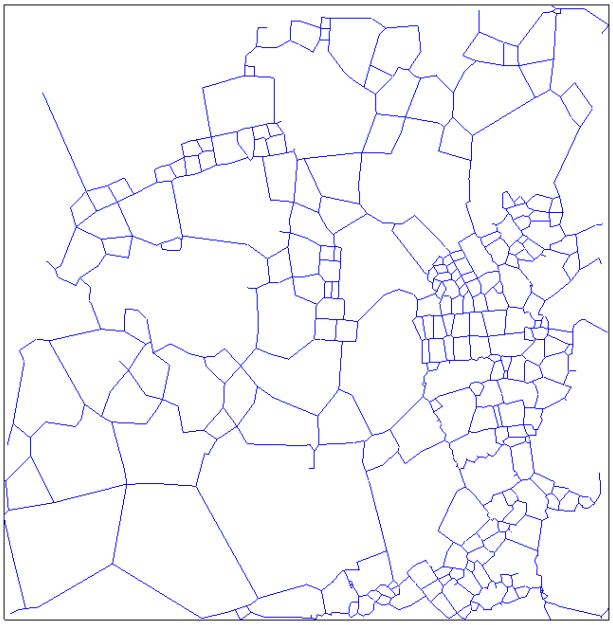

Figure 7: Relative Neighborhood Graph (RNG)

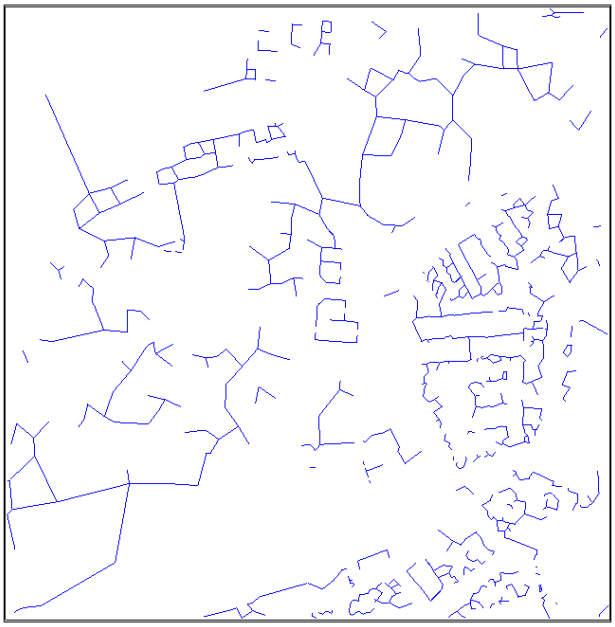

Figure 8: Result of clustering

5 CONCLUSION AND OTHER APPLICATIONS

The aim of this research is to find aggregated structures in spatial data sets similar to the ones defined in official maps or digital data sets. Another application is in the context of image interpretation, where the result of a supervised classification method is a collection of pixels with grayvalues representing different object classes. In order to delineate meaningful objects and their boundaries, region growing procedures have to be applied. If the pixels are however not coherent such procedures cannot be used. However, the problem can be tackled with the presented clustering approach, by transforming the pixels into the graph structure. Such an approach can e.g. then be used for the interpretation of satellite or aerial images in order to evaluate the growth rate of cities. In many regions of the world, this data is not documented nor planned - such techniques can help to quickly assess the relevant data in order to support city planning.

The clustering process presented can also be applied in a hierarchical manner: after clustering and derivation of cluster characteristics, new objects (i.e. the new clusters) are identified. These new objects, in turn, can be introduced in the clustering process again, after they are transformed into a node-edge graph structure, etc. In this way, a hierarchical, pyramid-like structure can be established. Thus the same cluster principle can be applied – however the neighborhood and similarity criteria are different and have to be derived dynamically from the given data. In the context of settlement aggregation, the following hierarchical structure can be derived: buildings $\rightarrow$ parts of the city $\rightarrow$ whole city $\rightarrow$ region.

REFERENCES

Anders, K.-H. & Sester, M. [1997], Methods of Data Base Interpretation - Applied to Model Generalization from Large to Medium Scale, *in:* W. Förstner & L. Plümer, eds, 'SMATI '97: Semantic Modelling for the Acquisition of Topographic Information from Images and Maps', Birkhäuser, pp. 89–103.

Ball, G. & Hall, D. [1965], 'ISODATA: a novel method of data analysis and pattern classification', Stanford Research Institute AD 699616.

Edelsbrunner, H. & Mücke, E. [1994], Three-dimensional Alpha Shapes, *in:* 'ACM Transactions on Graphics', Vol. 13, pp. 43–72.

Edelsbrunner, H., Kirkpatrick, D. & Seidel, R. [1983], On the shape of a set of points in the plane, *in:* 'IEEE Transactions on Information Theory', Vol. IT-29, pp. 551–559.

Ester, M., Kriegel, H.-P., Sander, J. & Xu, X. [1996], A Density-Based Algorithm for Discovering Clusters in Large Spatial Databases with Noise, *in:* 'Proceedings of 2nd. International Conference on Knowledge Discovery and Data Mining (KDD-96)'.

Gabriel, K. & Sokal, R. [1969], Anew statistical approach to geographic variation analysis, *in:* 'Systematic Zoology', Vol. 18, pp. 259–278.

Guha, S., Rastogi, R. & Shim, K. [1998], CURE: An efficient clustering algorithm for large databases, *in:* 'Proc. of 1998 ACM-SIGMOD International Conference on Management of Data'.

Guha, S., Rastogi, R. & Shim, K. [1999], ROCK: A robust clustering algorithm for categorical attributes, *in:* 'Proc. of the 15th International Conference on Data Engineering'.

Jain, A. & Dubes, R. [1988], *Algorithms for Clustering Data*, Prentice Hall.

Jaromczyk, J. & Toussaint, G. [1992], Relative neighborhood graphs and their relatives, *in:* 'Proceedings IEEE', Vol. 80, pp. 1502–1517.

Jarvis, R. [1977], Computing the shape hull of points in the plane, *in:* 'Proceedings of the IEEE Computing Society Conference on Pattern Recognition and Image Processing', pp. 231–241.

Jarvis, R. & Patrick, E. [1973], Clustering using a similarity measure based on shared near neighbours, *in:* 'IEEE Transactions on Computers', Vol. 22, pp. 1025–1034.

Karypis, G., Han, E.-H. S. & Kumar, V. [1999], CHAMELEON: A Hierarchical Clustering Algorithm Using Dynamical Modeling, To appear in the IEEE Computer or via internet at http://winter.cs.umn.edu/ karypis/publications/data-mining.html.

Kaufman, L. & Rousseeuw, P. [1990], *Finding Groups in Data: An Introduction to Cluster Analysis*, John Wiley & Sons.

Kirkpatrick, D. & Radke, J. [1985], A framework for computational morphology, *in:* G. Toussaint, ed., 'Computational Geometry', North-Holland, pp. 217–248.

MacQueen, J. [1967], Some methods for classification and analysis of multivariate observations, *in:* 'Proc. of the Fifth Berkeley Symposium on Mathematical Statistics and Probability', Vol. 1, pp. 281–297.

Ng, R. & Han, J. [1994], Efficient and effective clustering method for spatial data mining, *in:* 'Proc. of 1994 Int. Conf. on Very Large Data Bases (VLDB'94)', Santiago, Chile, pp. 144–155.

Plazanet, C., Bigolin, N. & Ruas, A. [1998], 'Experiments with Learning Techniques for Spatial Model Enrichment and Line Generalization', *GeoInformatica* **2**(4), 315–334.

Richardson, D. [1996], 'Automatic Processes in Database Building and Subsequent Automatic Abstractions', *Cartographica, Monograph 47* **33**(1), 41–54.

Ruas, A. & Lagrange, J. [1995], Data and knowledge modelling for generalization, *in:* J.-C. Müller, J.-P. Lagrange & R. Weibel, eds, 'GIS and Generalization - Methodology and Practice', Taylor & Francis, pp. 73–90.

Sester, M. [1999], 'Knowledge Acquisition for the Automatic Interpretation of Spatial Data', *Accepted for Publication in: International Journal of Geographical Information Science*.

Sester, M., Anders, K.-H. & Walter, V. [1998], 'Linking Objects of Different Spatial Data Sets by Integration and Aggregation', *GeoInformatica* **2**(4), 335–358.

Supowit, K. [1983], 'The relative neighborhood graph, with an application to minimum spanning trees', *J.Assoc.Comput.Mach.* **30**, 428–448.

Toussaint, G. [1980], The relative neighborhood graph of a finite planar set, *in:* 'Pattern Recognition', Vol. 12, pp. 261–268.

Toussaint, G. [1988], A graph-theoretical primal sketch, *in:* G. Toussaint, ed., 'Computational Morphology', North-Holland, pp. 229–260.

Toussaint, G. [1991], Some unsolved problems on proximity graphs, *in:* D. Dearholt & F. Harary, eds, 'Proceedings of the First Workshop on Proximity

Graphs. Memoranda in Computer and Cognitive Science MCCS-91-224', Computing research laboratory, New Mexico State University, La Cruces.

van Smaalen, J. [1996], 'Spatial Abstraction based on Hierarchical Re-Classification', *Cartographica, Monograph 47* **33**(1), 65–74.