Segmentation Based on Normalized Cuts for the Detection of Suburban Roads in Aerial Imagery

Anne Grote, Matthias Butenuth, Markus Gerke, Christian Heipke
Institute of Photogrammetry and GeoInformation
Leibniz Universität Hannover
Hannover, Germany
grote@ipi.uni-hannover.de

Abstract—**This paper deals with the segmentation of images of suburban scenes with the Normalized Cut algorithm. The segmentation results are intended to be used for the extraction of roads in order to assess existing road data. The similarity matrix necessary for the Normalized Cuts algorithm is built up using similarity criteria that are suitable for the separation of road segments and non-road segments. These criteria are edges, colour, hue and road surface colour derived with the help of the database information which is thus used as prior information to facilitate the segmentation and extraction. Segmentation is the main topic of this paper, but some hints on future work regarding the selection of road segments based on road colour are given. The results show that the approach is suitable for the segmentation in order to extract roads in suburban scenes.**

I. INTRODUCTION

GIS data become more and more important, and the need for high quality geospatial data increases. It is thus necessary to regularly check existing data to ensure that they are correct and up-to-date. The check against ground truth can be carried out with aerial images. Doing this manually is very time-consuming, therefore automation of the checking process is desirable. Roads are very important objects in databases, e.g. due to their widespread use in navigation systems, and they are subject to frequent change, especially in suburban areas: new roads are built and existing roads are relocated. An out-of-date database does not contain these roads at all or at least not correctly, or it contains vanished objects. In addition, there are roads which do not comply with the required positional or thematic accuracy. Keeping road databases up-to-date is very important, and in order to do this task automatically from images, proper road extraction algorithms are needed.

In the literature numerous approaches for road extraction exist, for example [1] and [2], and many of them work well, especially in open landscape. In urban and suburban areas the task is much more difficult, as is stated in [3]. The difficulties arise from the complex surroundings: buildings and trees next to the roads cast shadows, and they can cover parts of the road. More vehicles are to be expected than in rural areas. Straight road segments are more often interrupted by junctions and thus are shorter. One of the few approaches that were developed for built-up areas is the approach described in [4]. It is a very sophisticated system for the detection of road networks in inner-city areas that is based on grouping small entities step by step to lanes, carriageways and road networks. It works well in the test area but needs a multitude of parameters. The author states that additional information is essential for road extraction in urban areas due to the complexity of the surroundings. In his work he uses a digital surface model as additional information.

In this paper we present an approach for the extraction of roads in suburban areas using aerial imagery. In a later step the results are used for the verification of an existing road database, i.e. we aim to confirm or reject the correctness of road objects stored in the database. The database information is used as additional information to facilitate the extraction. For image segmentation – the first step towards road extraction – we use Normalized Cuts which is a graph-based method. Its advantage is the combination of several attributes in one similarity criterion which is used for segmentation. Afterwards, the segments are evaluated to differentiate road segments from non-road segments and to assess the quality of road information contained in the database. The evaluation and assessment stages are not covered in this paper.

II. NORMALIZED CUTS FOR IMAGE SEGMENTATION

In this section, the Normalized Cuts method, which is used for segmentation, is described in brief. Generally, the Normalized Cuts method is a method for the division of arbitrary undirected graphs on the basis of their weighted edges. The use of the method for image segmentation is described in [5]. The pixels are defined as nodes and connected among themselves by weighted edges. Theoretically, it is possible to connect one pixel to all other pixels, but in practice only some neighbouring pixels have weights different from zero. One edge weight is a measure for the similarity between two connected pixels. How this similarity is defined depends on the aim of the segmentation. For example, if the segments that are to be divided have distinctly different intensities, intensity should be used as similarity measure. Often, several similarity measures are combined, because the differentiation cannot be realised using only one criterion. The weight as the combination of all similarity measures has a value between 0 and 1. The graph representing the image is then cut into segments aiming at a large dissimilarity between different segments and at the same time a large similarity within the

1-4244-0712-5/07/$20.00 ©2007 IEEE.

segments. In order to achieve this goal, the graph is cut along the edges meeting the following condition:

$$Ncut(A,B) = \frac{link(A,B)}{link(A,V)} + \frac{link(A,B)}{link(B,V)} = \min \quad (1)$$

The graph is partitioned into two parts. *A* and *B* are the sets of nodes in the two separated subgraphs. *V* is the set of all nodes in the graph. *Link* is the sum of all weights connecting two sets of nodes:

$$link(P,Q) = \sum_{p \in P, q \in Q} w(p,q) \quad (2)$$

with *w(p,q)* as the edge weight between the individual nodes *p* and *q*. The weights for the pixel pairs are inserted in a symmetric similarity matrix (*W-matrix*) whose row and column dimensions equal the number of pixels. The minimum can be found by calculating the eigenvectors of a matrix derived from the *W-matrix*. The details of this calculation are described in [5]. The result is a set of discretised eigenvectors, as many as the desired number of segment, which has to be specified beforehand. The length of the eigenvectors corresponds to the number of pixels in the image. Each eigenvector indicates one segment.

III. Approach

The extraction of roads is the first step in the process of automatically assessing a road database. The extracted roads are later compared to the roads in the database with respect to existence and positional accuracy.

Because of disturbances like shadows and the inhomogeneous background, the line based road model frequently used in rural settings cannot be applied here. For this reason we use an area based approach. Difficulties arise on the one hand from the existence of disturbances, on the other hand from the fact that in an urban setting, many areas with similar colour and texture compared to roads can be found. Therefore, our road extraction approach does not rely on just one feature of roads, but it takes into account a combination of several features. In addition, we use the database information as prior knowledge.

A. Model

The road model (Fig. 1) that is used to select criteria for the similarity matrix and to evaluate the segments later is adapted from [6]. In this model one road segment consists of at least one lane which is represented in the image by a homogeneous area. One lane is bordered by road markings, represented in the image as bright lines, or the roadside, separated from the lane by an edge. The borders are parallel. Another type of road markings are symbols (e.g. zebra crossings or arrows) which are not used explicitly yet.

The appearance of road objects in aerial imagery can be disturbed by context objects, resulting in the need for a context model. The context model shown in Fig. 2 is also adapted from [6]. Context objects that can have an effect on the appearance of the road are buildings and high vegetation, such as trees, which can stand close to the road. They directly occlude roads in aerial imagery or may cast shadows. Furthermore, vehicles may hide parts of the road.

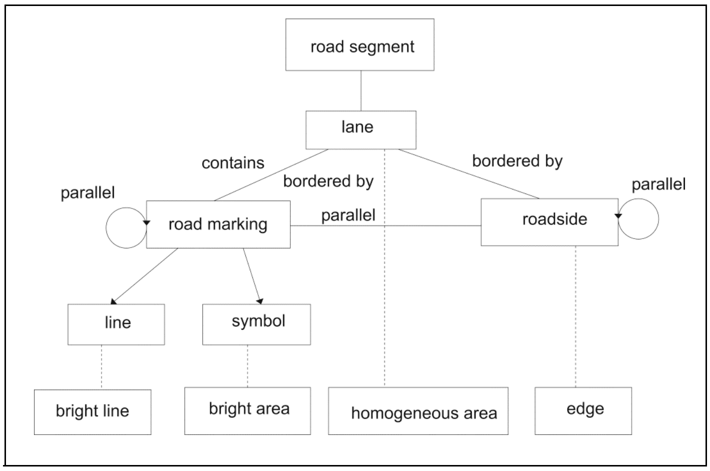

Figure 1. Road model (adapted from [6]).

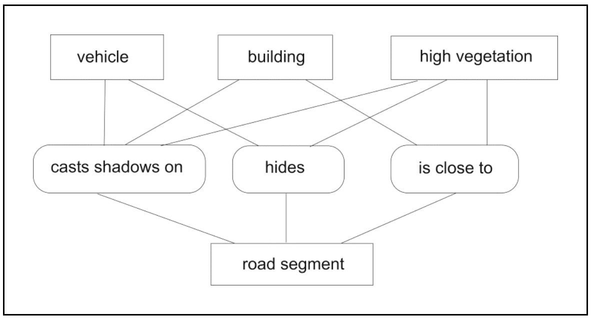

Figure 2. Context model (adapted from [6]).

B. Strategy

The verification of a given road database is the final goal of our research. We use this road data as prior information for the segmentation, i.e. to define areas of interest around the road centre line taken from the database. Its width is set to approximately three times the road width taken from the database, allowing for both road and context objects to be identified. This is a reasonable procedure, because during verification only the given data is to be checked – roads outside this buffer are not of interest.

The image is segmented within the area of interest using Normalized Cuts, as described in detail in subsection C. Afterwards, the segments have to be evaluated and grouped to road segments and non-road segments. Road segments are merged to form a road which is then used to assess the database road. This last step is not covered in this paper.

C. Segmentation

For road extraction the image first has to be segmented which is done here using Normalized Cuts (see section II). Since the segmentation using Normalized Cuts is computationally rather expensive, the image is not segmented as a whole but each road in the database is examined individually, for which purpose an area of interest is defined

1-4244-0712-5/07/$20.00 ©2007 IEEE.

around the position of the database road. The segmentation takes place in this area of interest.

For segmentation with Normalized Cuts, the similarities between pixel pairs have to be defined. The selection of similarity measures that are combined to generate the weight of each pixel connection is based on the road and context model.

As roads are separated from their surroundings by edges, the first similarity criterion is the existence of edges between two pixels. For this purpose, the Laplace-of-Gaussian operator is applied to the image yielding an image of the second derivative, where edges are indicated by sign changes (zero-crossings). We then check whether a sign change occurs along the connecting line between the two pixels. If there is a sign change, the edge amplitude at this point is taken from an edge image calculated by the Canny operator and inserted in the following equation, according to the way described in [5], to obtain the first part of the similarity measure:

$$w_{edge} = e^{\frac{-f^2}{2\sigma^2}} . \tag{3}$$

Here, f is the edge intensity and σ is ten percent of the range of the edge intensity, following the example of [5].

The second criterion is colour because the colour of the road surface normally does not change strongly. A measure for the colour similarity of two pixels is the length of the distance vector of the two colour vectors. The second part of the similarity measure is then calculated as in (3), with f as the length of the difference vector and σ defined respectively.

Both criteria are then multiplied:

$$w = w_{edge} \cdot w_{colour} . \tag{4}$$

A third criterion is hue because in this way the problems related to shadow can be at least partly solved: an object keeps approximately the same hue in a colour image even in the parts darkened by shadows [7]. Therefore, the weight is reduced if the two pixels have a hue difference that is greater than a threshold.

To further support the segmentation the database information is used to obtain colour information about the roads: from the position of the database road axis, the average colour values of the road are calculated. Then we check for every pair of pixels if both pixels lie within a tolerance interval of the average values defined by the standard deviation. If both pixels lie inside or outside this interval, the weight is increased.

After the *W-matrix* is built out of these similarity criteria, the Normalized Cuts algorithm is carried out. The number of segments has to be specified before the algorithm is started. It must be large enough in order to prevent merging of road and non-road segments.

IV. RESULTS

For the following examples CIR aerial images with a resolution of 0.1 m per pixel were used. The images are subsets from a suburban scene in Grangemouth, Scotland. As additional information, we have a digital surface model with a horizontal resolution of 0.2 m and a height resolution of 0.1 m.

Fig. 3 and 4 show examples of segmentation results with Normalized Cuts. The database road is indicated by a thick black line. The images are separated into parts to increase performance. Each part is divided into 20 segments by the Normalized Cuts algorithm. The results show that the segmentation in general has succeeded: most segments cover either only a road area or only a non-road area. Exceptions are places where trees cast shadows on the road, showing that in these images, the hue feature has not much influence, and places where the contrast between the road and its surroundings is low. The algorithm mostly ignores small surface changes and yields relatively smooth segments.

The benefit of using database information for the determination of road colour is demonstrated by a comparison between Fig. 4 and Fig. 5. The segments displayed in Fig. 5 are computed without taking the database information into account for the road colour, but otherwise with the same parameters. In Fig. 5 several road segments also contain pavement areas, for example in the lower right of the second part. In Fig. 4 this only happens when the contrast is low, in the right parts of the image. In general, the results are similar, but the database information can improve the result.

Fig. 6 shows a segmentation result in an area with a false database road. Here the database information is used in the same manner as in Fig. 3 and 4. This example shows that the use of wrong database information does not cause a wrong segmentation. The segments are defined by the image content, they do not line up along the false road. So, the database information can improve the segmentation results if it is correct, but it does not corrupt the results if it is wrong.

V. CONCLUSIONS AND OUTLOOK

The examples that were displayed in this paper show the general usability of the approach for the detection of roads in connection with a road database. The Normalized Cuts algorithm is suitable for the segmentation step. One advantage of this method is the possibility to combine several different segmentation criteria in one step. Another advantage is that both local and global characteristics of the image are taken into account. Local characteristics are incorporated into the similarity matrix which considers the similarity of pixels in a close neighbourhood. Global characteristics come into play when the best cut is computed: a global minimum criterion must be met. We believe that this combination of local and global aspects is a very important characteristic of the Normalized Cuts algorithm.

The next steps of our work are the selection and grouping of road segments. In order to accomplish them, the segments need to be evaluated with respect to their shape and colour. The NDVI is used to exclude areas with vegetation and the digital surface model is used to exclude high areas and to explain disturbances due to shadows. A first examination has been carried out where colour is used to extract some preliminary road segments. All segments whose mean colour lies inside the standard deviation interval of the average road colour

1-4244-0712-5/07/$20.00 ©2007 IEEE.

determined in the segmentation step are selected. Fig. 7 displays the segments that are selected from the segmentation result in Fig. 4 using this constraint. About half of the real road segments were selected, only one selected element does not belong to the road but was selected because it has the same surface. From this start, it is possible to merge adjacent road segments and evaluate their shape and direction. Following the road direction, new segments can be added using shape and colour criteria. When no segment is found, the reason is searched for with the help of the context model. Dark, irregularly shaped segments, for example, can be interpreted as shadows when there is a high object near by.

In Fig. 8, the result of the segment selection for the false database road (from Fig. 6) is shown. The false road crosses buildings and trees as is to be expected in a suburban scene. For this reason, height data can be used to rule out many segments. High objects are depicted in white in Fig. 8. The remaining selected segments are shown with white borders as in the previous examples. Two of them can be ruled out due to high NDVI, irregular shape and missing neighbours. The other segments actually belong to a road. So, in this image, a road would be found, but in the following assessment step the false database road would be rejected because it has another direction.

ACKNOWLEDGMENT

This project is funded by the DFG (German Research Foundation). The calculations are made with a C++ program adapted from a MATLAB program written by Timothée Cour, Stella Yu and Jianbo Shi. Their program can be found on http://www.seas.upenn.edu/~timothee/software_ncut/software.html (last checked Jan 2007).

REFERENCES

[1] E.P. Baltsavias, "Object extraction and revision by image analysis using existing geodata and knowledge: current status and steps towards operational systems", *ISPRS Journal of Photogrammetry and Remote Sensing*, vol. 59, issue 1-2, pp. 73-84, 2004.

[2] C. Wiedemann and H. Ebner, "Automatic completion and evaluation of road networks", *International Archives of Photogrammetry and Remote Sensing*, Amsterdam, The Netherlands, vol. 33, part 3B, pp. 979-986, 2000.

[3] H. Mayer, S. Hinz, U. Bacher, and E. Baltsavias, "A test of automatic road extraction approaches", *International Archives of photogrammetry, remote sensing and spatial information sciences*, vol. 36, part 3, pp. 209-214, 2006.

[4] Hinz, S., "Automatische Extraktion urbaner Straßennetze aus Luftbildern", *Dissertation, TU München, DGK Reihe C, Nr. 580: Verlag der Bayerischen Akademie der Wissenschaften*, 2004, 130 p.

[5] J. Shi and J. Malik, "Normalized cuts and image segmentation", *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 22(8), pp. 888-905, 2000.

[6] S. Hinz and A. Baumgartner, "Urban road net extraction integrating internal evaluation models", *International Archives of Photogrammetry, Remote Sensing and Spatial Information Sciences*, vol. 34, part 3/A, pp. 163-168, 2002.

[7] F. Perez and C. Koch, "Toward color image segmentation in analog VLSI: algorithm and hardware", *International Journal of Computer Vision*, vol. 12, no. 1, pp. 17-42, Feb. 1994.

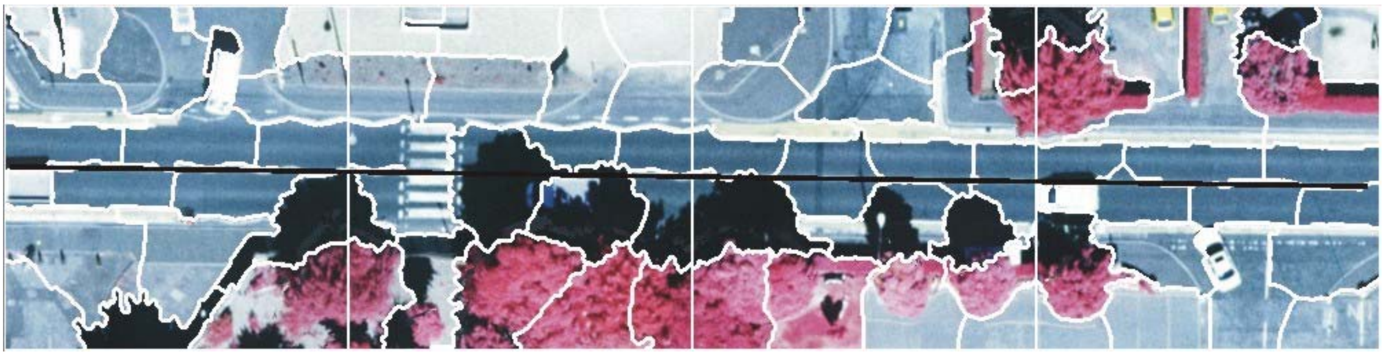

Figure 3. Image segmented with Normalized Cuts, first example.

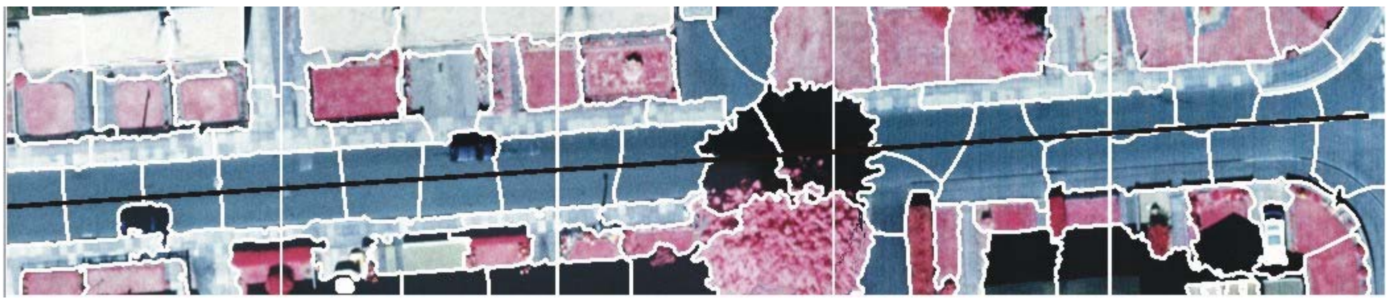

Figure 4. Image segmented with Normalized Cuts, second example.

1-4244-0712-5/07/$20.00 ©2007 IEEE.

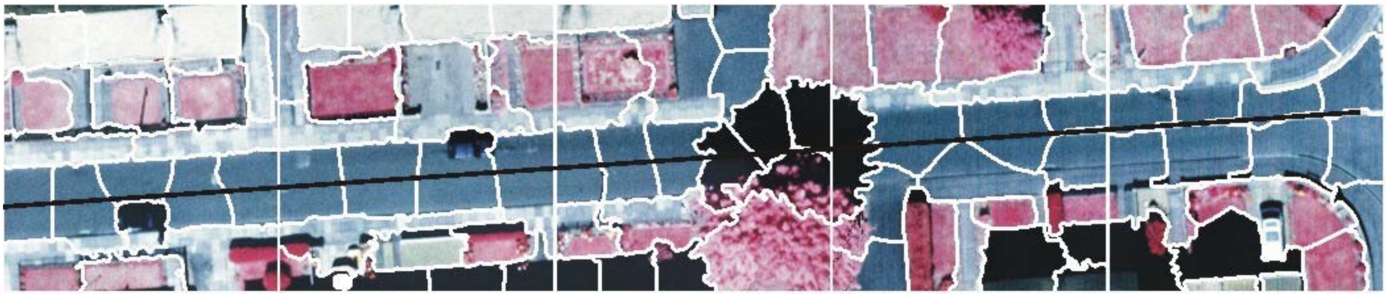

Figure 5. Segmentation without using road colour information derived from the database road.

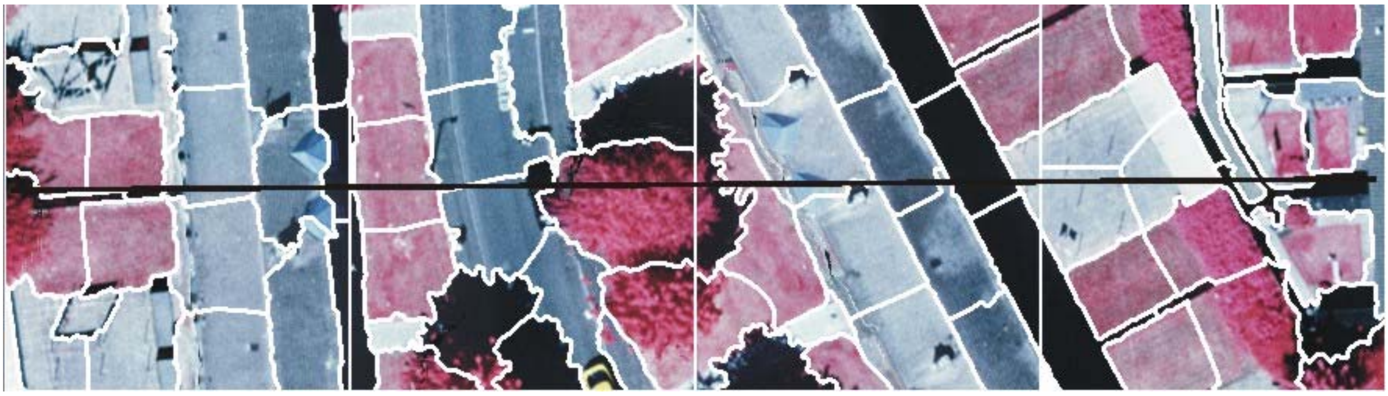

Figure 6. Example for a false database road.

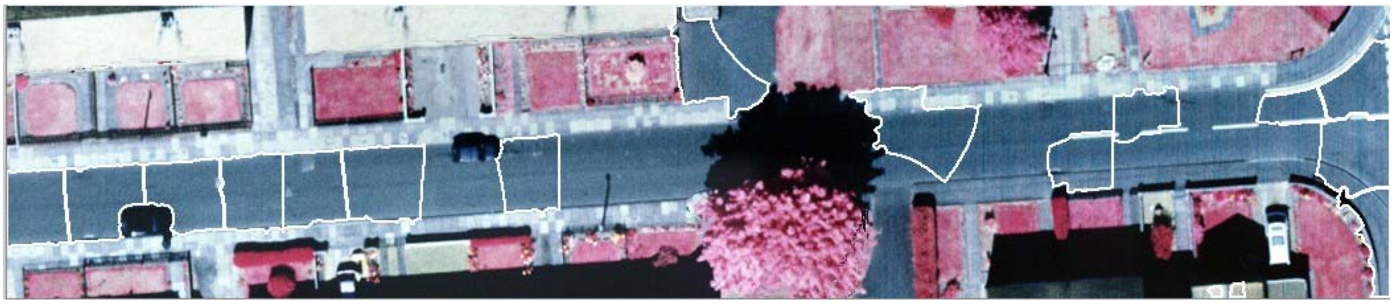

Figure 7. Selected road segments from first example

Figure 8. Selected segments and high objects in image with false database road.