Road extraction in suburban areas by region-based road subgraph extraction and evaluation

Anne Grote, Christian Heipke, Franz Rottensteiner, Hannes Meyer
Institute of Photogrammetry and GeoInformation
Leibniz Universität Hannover
Hannover, Germany
grote@ipi.uni-hannover.de

Abstract—**In this paper, a road extraction approach for suburban areas from high resolution CIR images is presented. The approach is region-based: the image is first segmented using the normalized cuts algorithm, then the initial segments are grouped to form segments, and road parts are extracted from these segments. Ideally roads in the image correspond to only one extracted road part, but they are often covered by several road parts with gaps between them. In order to combine these road parts, neighbouring road parts are connected to a road subgraph if there is evidence that they belong to the same road, such as similar direction and smooth continuation. This process allows several branches in the subgraph which is why another step follows to evaluate the subgraphs and divide them at gaps which show weak connections. The subgraph evaluation step is the focus of this paper. Linear programming is used for the subgraph evaluation after gap weights are determined. Two ways of determining gap weights are discussed, one using criteria which describe the properties of the road parts and their interrelations, and one using context objects (vehicles, trees, vegetation) in the gaps. The determination of the gap weights and the division of the road subgraphs is shown with an example.**

I. INTRODUCTION

Roads are an important part of infrastructure, especially in an urban context. Road data are needed in many applications, and the information must be up-to-date and correct if it is to be useful. In order to keep road databases up-to-date and accurate, aerial and satellite images are often used to gather information about changes, because the images are relatively cheap and often much more accessible compared to field measurements. In order to further reduce the costs and the time required for map updating, it is desirable to use automatic procedures for the extraction of roads from the images. Today, especially in urban areas roads are to a large degree still extracted manually. This is due to the considerable complexity of urban environments compared to open landscapes, for which road extraction algorithms that are reasonably reliable already exist, e.g. [1], [2].

There are many different approaches for road extraction from optical imagery, and in recent years the number of those that deal with urban areas has increased. For urban areas, high resolution images (1 m and better) are used almost exclusively. The only exceptions are approaches where only main roads are extracted, e.g. [3]. Most approaches are based on images with 1 m resolution, but there are also some examples where images with 0.1-0.2 m resolution are used, most notably [4], [5]. This is also one of the rare examples for road extraction in urban areas where context objects, namely buildings, vehicles, and shadow areas, are explicitly used. This can prevent gaps in the road network that are caused by the influence of these objects.

The high complexity of urban and suburban areas makes road extraction from greyscale aerial images without further information difficult because many different structures in urban areas have an appearance similar to that of roads. Therefore, most approaches use additional information, for example colour [6], [7], digital surface models [4] or both [8]. Information about the position of roads from an existing road database is also included sometimes [9]. Very few approaches use only greyscale images, for example [10], which is a tracking approach that can use automatic as well as manual seeds.

Most road extraction algorithms are either line-based or region-based. Line-based approaches for high resolution images mostly extract edges of roads and group them to form road lanes [4]. In region-based algorithms, colour images are often classified and the road regions found by classification are refined by morphological operators and/or selected according to shape criteria [6], [9].

Our goal is the extraction of roads in suburban areas. We use a region-based approach in which the image is first segmented and then road parts are extracted from the segments. These road parts are assembled into road subgraphs which can contain different branches. The branches represent different hypotheses for the course of the roads, from which the correct one has to be picked in order to obtain a consistent road network in the next step. This paper focuses on the evaluation of the road subgraphs to find the most likely course of the road from different hypotheses, using either relations between road parts or relations to context objects. In section II of this paper the approach is explained. The segmentation, road part extraction and subgraph generation, which are explained in detail in [11], [12], are only reviewed briefly, the road subgraph evaluation is discussed in more detail. In section III some results of the road subgraph evaluation are presented. Section IV gives conclusions and pointers for further work.

978-1-4244-3461-9/09/$25.00 ©2009 IEEE

II. Approach

A. Overview

The goal of this approach is the extraction of roads from high resolution aerial images in suburban areas. We use colour infrared images with a resolution of approximately 10 cm on the ground.

The approach consists of three steps, namely segmentation, road part extraction and subgraph generation (see [11] and [12] for more details). In the segmentation step, the image is first divided into many small segments, which are then grouped into larger segments having meaningful shapes. From the grouped segments, potential road parts are extracted using shape criteria. The road parts are then assembled to road subgraphs if they potentially belong to the same road; junctions are not considered. This step allows several branches to be present in one subgraph. In the next step, these ambiguities are resolved by optimizing the graph in a way that finds the best possibility for the course of the road without branches.

The use of the terms *road part* and *road subgraph* is explained in Fig. 1. An extracted road part is a segment which was classified as a road. It can correspond to a whole actual road between two junctions or only a part of the road, or it can be a false positive. A road subgraph consists of several assembled road parts. Each subgraph extends only as far as there are road parts to be found in a more or less straight continuation; in this way, each subgraph usually represents only one road. In the subgraph, each road part has two *nodes* which are connected via an edge called *road edge*. Each node can also maintain connections to nodes of other road parts via edges called *gap edges*. If more than one such connection exists, the node has several *branches*. The term *subgraph* is used in order to indicate that it does not represent a complete, interconnected road network.

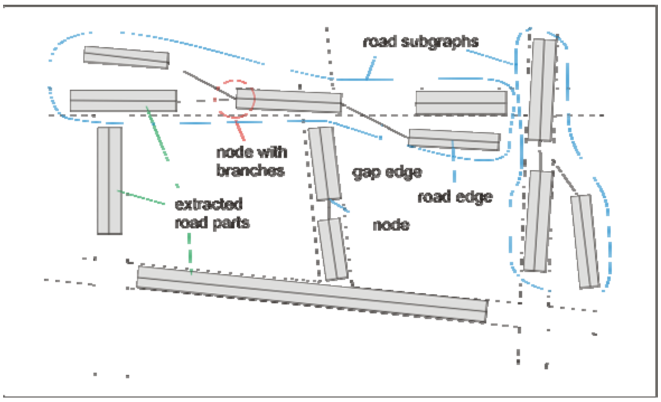

Figure 1. Definition of terms. Dashed line: real road network; grey rectangles: extracted road parts; black lines between road parts: edges of road subgraphs.

B. Segmentation, Road Part Extraction and Road Subgraphs

The first stage of the road extraction is the segmentation of the image, which is carried out in two steps, namely initial segmentation and grouping. The goal of the initial segmentation is to divide the image into small regions whose borders coincide with the road borders as completely as possible. The normalized cuts algorithm [13] is used for this initial segmentation, in which connections between pixels are weighted according to their similarities. The similarities of pixel pairs are determined using colour and edge criteria.

As the normalized cuts algorithm results in a considerable oversegmentation in order to preserve most road borders, the initial segments are grouped. The grouping is done iteratively using colour and edge criteria, this time considering the properties of the regions as opposed to those of the pixels. Segments with irregular forms that cover roads across junctions can occur in this step. Therefore, the skeletons of the segments are examined and if they have several long branches, the segments are split.

Road parts are extracted from the grouped segments in the next step. For the extraction, geometric and radiometric criteria are used. The geometric criteria are elongation, width constancy and difference to average road width. As radiometric criteria, the NDVI (normalized difference vegetation index) and the standard deviation of colour are used, dark areas are excluded because shadow areas often have similar geometric properties to road parts. The elongation, width constancy, compliance with average road width and the NDVI are used to determine a quality measure for each road part. For each road part an evaluation score is computed, further details of the road part extraction are explained in [12].

In many cases, a road is not completely covered by one road part but by several different road parts because disturbances in the appearance of the road interfered with the extraction. Therefore, road parts that could belong to the same road are assembled into road subgraphs (Fig. 1) by checking if the road parts have neighbours to which they can be connected. The subgraphs are assembled iteratively, starting with the road part with the best evaluation result from the road extraction. The criteria used to decide whether two road parts belong to the same road are distance between segments, direction difference and continuation smoothness. The reference points for the measurement of the direction difference and the continuation smoothness are the intersection points between the center line and the road part borders. The continuation smoothness is measured by calculating the direction differences between the directions of the road parts to the direction of their connection (Fig. 2). The continuation smoothness is high if both smoothness angles are small.

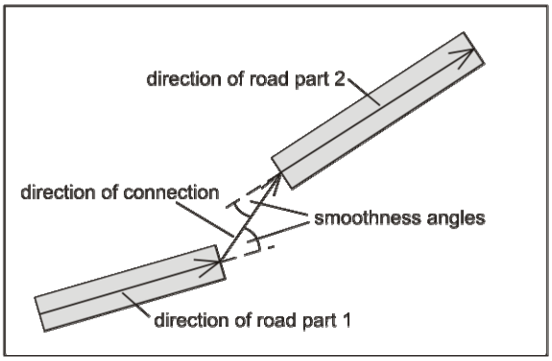

Figure 2. Continuation smoothness

If the distance between the segments is short, the continuation smoothness criterion is disregarded because at close distances the angles depend too much on the exact positions from which the angles are measured.

Two road parts are linked if empirically determined thresholds for the distance, the direction difference and the continuation smoothness are met. The distance and the direction difference must be low and the continuation smoothness must be high; all three conditions must be fulfilled for the road parts to be linked. One road part can be attached to more than one other road part in the same direction, such that branches in the subgraphs can occur (Fig. 1). The search for neighbouring road parts continues until no more road parts can be added. Then, the search is resumed with the road part which has the best evaluation result among the remaining road parts until all road parts have been examined.

C. Subgraph Evaluation by Linear Programming

As described in the previous section, the road subgraphs can contain branches. However, in most cases these branches do not represent actual branches in the road network but rather indicate false extractions of road parts that are nearly parallel to the real road. Therefore, road subgraphs containing branches are treated as including different hypotheses for the course of the road. In this step, the best hypotheses are searched for via the formulation and solution of a linear programming problem.

In linear programming a linear function (objective function) whose variables are subject to linear constraints is maximized or minimized [14]. The constraints define a set of feasible vectors; the vector for which the constraint set is maximal or minimal is the optimal solution for the problem.

Linear programming can be used when the variables of the linear function to be optimized are restricted by hard constraints, which can be equations or inequalities. The constraints in our case are inequalities which result from the condition that no node of a subgraph should be connected to more than one gap edge after the optimization.

The problem to be solved is finding the best partition of a road subgraph. Each gap edge has a weight which reflects the plausibility that the two road parts belong to the same road (see below). In order to keep good gap edges, the sum of the edge weights should be maximized for the edges which are kept. Thus, the objective function is

$$w_1 x_1 + ... + w_n x_n \rightarrow \max \qquad (1)$$

with n as the number of gap edges, $w_1 \ldots w_n$ as the respective weights for the gap edges and $x_1 \ldots x_n$ as the variables that indicate whether the respective edge should be kept or discarded. A value of 1 indicates that the edge is kept; a value of 0 indicates that it is discarded. These values are determined by solving the maximization under the constraints that each node i can only be associated to one gap edge:

$$\sum_{j \in E_i} x_j \leq 1 \cdot \qquad (2)$$

Here, E_i is the set of gap edges belonging to node i. The optimization is carried out using the simplex method [14].

The edge weights are determined using the following criteria:

- distance: a shorter distance between the two connected road parts gives a higher edge weight;
- road part quality: the sum of the quality measures of both road parts from the extraction, a higher value gives a higher edge weight;
- colour: a smaller difference between the mean colour values of both road parts gives a higher edge weight;
- width: a smaller width difference between both road parts gives a higher edge weight;
- continuation smoothness: smaller smoothness angles (cf. subsection B) give a higher edge weight;
- direction: a smaller direction difference between both road parts gives a higher edge weight.

The weights for the different criteria are obtained after calculating all criteria by mapping the respective values linearly onto an interval between 0 and 1. For example, the maximum possible distance between two connected road parts is equivalent to a distance weight of 0, and a distance of 0 is equivalent to a distance weight of 1. The other weights are obtained accordingly. All weights are multiplied to obtain the total weight for one edge. The edge weights that belong to the same subgraph are normalized such that their sum equals 1.

D. Additional Use of Context Objects

Context objects can be used to assist the determination of the gap weights. Some experiments have been carried out using an evaluation of the context objects for the determination of gap weights [15].

For the evaluation of a gap between two road parts the context objects vehicle, tree, shadow, vegetated area and asphalt area are used. Asphalt areas are areas that were not extracted as road parts but have the radiometric properties of roads. They are treated as context objects that support a road hypothesis. The context objects are extracted automatically as follows:

- Vehicles are extracted by filtering small dark or bright regions for rectangles with vehicle size or smaller; regions that represent hood, roof and rear are composed if they together form a vehicle.
- Trees are extracted as regions with high NDVI and associated shadows. Shadow areas are found by extracting dark areas.
- Vegetated areas are extracted as regions with high NDVI that are not classified as trees.
- Asphalt areas are extracted as areas with the average grey values of asphalt.

Details about the extraction of context objects can be found in [15].

To evaluate the gap, a road part hypothesis with the average width of the connected road parts is assumed in the gap. Two aspects of context objects are considered:

1. how much the existent context objects support or contradict the road part hypothesis in the gap, and
2. how much the existent context objects hinder the algorithm for the road extraction by occlusion.

For the first aspect the different relations between context objects and the road part hypothesis in the gap are classified into different categories. The relation categories for vehicles are *vehicle parallel on road*, *vehicle perpendicular to road*, *vehicle parallel next to road* and *vehicle perpendicular next to road*. A vehicle is counted as parallel if its direction differs less than 45° from the road direction; else it is counted as perpendicular. Parallel vehicles give higher evidence to support the road hypothesis than perpendicular vehicles; a parallel vehicle on the road gives higher evidence than one next to the road. The relation categories for trees are *tree on road*, *tree next to road* and *row of trees parallel next to road*. Trees on the road give evidence against a road hypothesis, single trees next to the road hypothesis gives a weak support which is stronger for a row of trees. The relation categories for vegetated areas are *vegetated area on road* and *vegetated area next to road*. Vegetated areas next to a road hypothesis give weak evidence to support the road hypothesis; if they are found on a road hypothesis, they give strong evidence against the hypothesis. The relation categories for asphalt areas are *asphalt area on road* and *asphalt area next to road*. Both give evidence to support a road hypothesis.

All context objects are grouped into these categories (shadows are only used in conjunction with the tree extraction). A relation value between -0.5 and 0.5 is assigned to each relation category; negative values indicate contradicting evidence while positive values indicate supporting evidence. These values are derived from the observed frequency of the respective relation in the global context of suburban areas and a weight for the importance of the relation in contradicting or supporting road hypotheses. If two or more relations of one category appear, the value for the second appearance is divided by two, the third by three and so on. The total evaluation for the first aspect is the sum of all single evaluations for the relations. If the relation *tree on road* or *vegetated area on road* appears, the road hypothesis receives a high negative value.

For the second aspect the occlusion of the road hypothesis by context objects is analysed. An occlusion can cause the road extraction algorithm to fail to extract the road part. So a high degree of occlusion by the context objects vehicle, tree and shadow supports a road hypothesis more than a low degree of occlusion. The estimation for the second aspect is the percentage of the area covered by the context objects in the road part hypothesis.

The final value of the context evaluation for the gap is the sum of the values from the two aspects.

III. RESULTS

Experiments were made on subsets of a colour infrared orthoimage with 10 cm resolution depicting a suburban scene from Grangemouth, Scotland.

Segmentation, grouping, road extraction and road subgraph formation were conducted as described in section II.B and [11, 12]. Fig. 3 shows a result on one subset after road subgraphs are assembled.

Figure 3. Road subgraphs. Different colours show different subgraphs.

At the bottom of the image, a road subgraph with branches can be seen: the left road part is connected to both the center road part and the bottom road part. The right road part is not connected to the bottom road part because the continuation smoothness is not good enough. This subgraph is now examined further with the linear programming optimization in order to remove the branching. Fig. 4 shows the subgraph with numbered gap edges, and Fig. 5 shows the subgraph with numbered nodes; the numbers will be used in the following description.

Figure 4. Road subgraph with numbered edges.

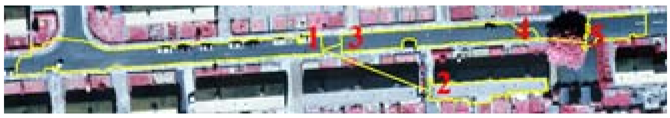

Figure 5. Road subgraph with numbered nodes.

The total weights for each edge are calculated and normalized (cf. II.C); they are:

$$w_1 = 0.03$$
$$w_2 = 0.85$$
$$w_3 = 0.11.$$

978-1-4244-3461-9/09/$25.00 ©2009 IEEE

The weight for edge 1 is low because of a high distance between the road parts, low continuation smoothness and high difference in colour. Edge 2 has the highest weight because of the short distance, small colour difference and good quality measures for both road parts. At edge 3 the distance is longer and the quality measure for the right road part is not as good.

The objective function for the linear program with these weights is

$$w_1\, x_1 + w_2\, x_2 + w_3\, x_3 \rightarrow \max.$$

The constraints indicate which nodes are connected to which edges (the numbers before the inequalities refer to the nodes):

$$\begin{array}{llll} 1: & x_1 + x_2 & & \leq 1 \\ 2: & x_1 & & \leq 1 \\ 3: & \quad x_2 & & \leq 1 \\ 4: & & x_3 & \leq 1 \\ 5: & & x_3 & \leq 1 \end{array}$$

The linear program is then solved and yields the following results:

$$x_1 = 0,\;\; x_2 = 1,\;\; x_3 = 1.$$

Thus, the result of the optimization is that edge 1 should be removed, leaving one road string consisting of three road parts and one separate road part (which is actually a false positive). Fig. 6 shows the road subgraph with the separated edge highlighted.

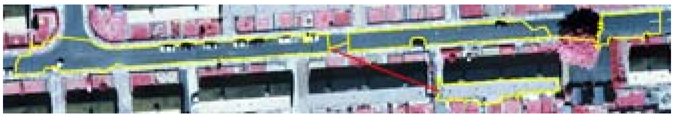

Figure 6. Road subgraph with separated edge.

As described in subsection II.D, context objects can be used to evaluate the gaps. An example for this is shown with the same road subgraph as above. The extracted context objects are shown in Fig. 7.

Figure 7. Context objects. Blue: vehicles, light green: trees, yellow: tree shadow, dark green: vegetated area.

For edge 1 (numbering as shown in Fig. 4), two vegetated areas are found next to the road hypothesis as well as asphalt (not displayed) near the branching node. The context evaluation result for this edge is c_1 = 0.606. The evaluation suffers from missing one vegetated area in the extracted context objects such that the evaluation result is relatively high for a false road hypothesis.

The context objects for edge 2 also are vegetated areas and asphalt area. The evaluation result is c_2 = 0.567. The gap here is quite short, which does not allow for many context objects, so the evaluation result is relatively low.

Edge 3 represents a gap which is caused by a tree standing next to the road. Consequently, the most prominent context object is the tree with its shadow. Vegetated areas and asphalt areas are also found. The evaluation result is c_3 = 0.659. The largest contribution to this value comes from the occlusion by the tree shadow.

If the evaluation results from the context objects alone are used as gap weights for the linear program, edge 2 is removed, which is not a desirable result. The reasons for this result are, as stated, the missing vegetated area at edge 1 and the few context objects at edge 2. On the other hand, the gap weight for edge 2 according to the criteria of subsection II.C is the highest of all three edges. This indicates that context objects should not be used alone without other criteria, especially the length of the gap.

IV. Conclusions

In this paper, an approach for the extraction of roads in suburban areas was presented, with the focus on resolving competing road hypotheses in a road subgraph. The task was formulated as a linear programming problem, and the application of the linear program was shown with one example of a road graph with one branch.

For the determination of the gap weights, two approaches were shown: the first using several criteria concerning properties of the road parts and their relations to each other (section II.C), the second using context objects in the gap (section II.D). In the example used here the first approach gave a better result, but this should not be generalized without further investigation. It is also planned to combine the context object evaluation with the other criteria as, for example, the context object evaluation for edge 2 (see section III) shows that context should not be used alone without regard for the length of the gap.

Two further ways for improving the determination of the gap weights are planned to be examined. The first is to improve combination of the criteria described in subsection II.D, especially the relative importance of the criteria and their relations to each other. The second is to include a digital surface model which allows to extract buildings automatically and use them as additional context objects. Preliminary experiments with manually extracted buildings show that this improves the context object evaluation significantly. A digital surface model can also be of use in the previous steps, especially the road part extraction.

The next steps also include the formation of a road network by searching for junction hypotheses between road strings and removing isolated (mainly false) road parts.

Acknowledgment

This project is funded by the DFG (German Research Foundation). The calculation of the normalized cuts was made with a C++ program partly adapted from a MATLAB program written by Timothée Cour, Stella Yu and Jianbo Shi. Their program can be found at http://www.seas.upenn.edu/~timothee/software_ncut/software.html (last checked March

2009). The calculation of the linear program was made with the MILP solver lp_solve (http://lpsolve.sourceforge.net/5.5/, last checked March 2009).

REFERENCES

[1] C. Zhang, “Towards an operational system for automated updating of road databases by integration of imagery and geodata”, *ISPRS Journal of Photogrammetry and Remote Sensing*, vol. 58, no. 3-4, 2004, pp. 166-186.

[2] H. Mayer, S. Hinz, U. Bacher and E. Baltsavias, „A test of automatic road extraction approaches“, *International Archives of Photogrammetry and Remote Sensing*, vol. 36, part 3, 2006, pp. 209-214.

[3] X. Hu and V. Tao, “Automatic extraction of main road centerlines from high resolution satellite imagery using hierarchical grouping”, *Photogrammetric Engineering & Remote Sensing*, vol. 73, no. 9, 2007, pp. 1049-1056.

[4] S. Hinz, “Automatic road extraction in urban scenes – and beyond”, *International Archives of Photogrammetry and Remote Sensing*, vol. 35, part B3, 2004, pp. 349-355.

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

[6] Q. Zhang and I. Couloigner, “Automated road network extraction from high resolution multi-spectral imagery”, *Proceedings of ASPRS Annual Conference*, Reno, Nevada, May 2006, 10 p., CD.

[7] P. Doucette, P. Agouris and A. Stefanidis, “Autometed road extraction from high resolution multispectral imagery”, *Photogrammetric Engineering & Remote Sensing*, vol. 70, no. 12, 2004, pp. 1405-1416.

[8] X. Hu, C. V. Tao and Y. Hu, “Automatic road extraction from dense urban area by integrated processing of high resolution imagery and LIDAR data”, *International Archives of Photogrammetry and Remote Sensing*, vol. 35, part B3, 2004, pp. 288-292.

[9] J.B. Mena and J.A.Malpica, “An automatic method for road extraction in rural and semi-urban areas starting from high resolution satellite imagery”, *Pattern Recognition Letters*, vol. 26, no. 9, 2005, pp. 1201-1220.

[10] J. Hu, A. Razdan, J.C. Femiani, M. Cui and P. Wonka, “Road network extraction and intersection detection from aerial images by tracking road footprints”, *IEEE Transactions on Geoscience and Remote Sensing*, vol. 45, no. 12, 2007, pp. 4144-4157.

[11] A. Grote and C. Heipke, “Road extraction for the update of road databases in suburban areas”, *International Archives of Photogrammetry and Remote Sensing*, vol. 37, part B3b, 2008, pp. 563-568.

[12] A. Grote, M. Butenuth and C. Heipke, “Road extraction in suburban areas based on normalized cuts”, *International Archives of Photogrammetry and Remote Sensing*, vol. 36, part 3/W49A, 2007, pp. 51-56.

[13] J. Shi and J. Malik, “Normalized cuts and image segmentation”, *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 22, no. 8, 2000, pp. 888-905.

[14] T. S. Ferguson, “Linear programming: A concise introduction”, *UCLA [online]*. Available: http://www.math.ucla.edu/~tom/LP.pdf [last accessed 2009, February 5].

[15] H. Meyer, “Verwendung von Kontextobjekten für die Extraktion von Straßen in Vorstadtgebieten aus Luftbildern“, *Master’s Thesis, Leibniz Universität Hannover*, 2009.

978-1-4244-3461-9/09/$25.00 ©2009 IEEE