

ASSESSING THE IMPACT OF DIGITAL SURFACE MODELS ON ROAD EXTRACTION IN SUBURBAN AREAS BY REGION-BASED ROAD SUBGRAPH EXTRACTION

Anne Grote, Franz Rottensteiner

Institute of Photogrammetry and GeoInformation, Leibniz Universität Hannover, 30167 Hannover, Germany
(grote, rottensteiner)@ipi.uni-hannover.de

Commission III, WG III/4

KEY WORDS: High resolution, Aerial, Urban, Automation, Extraction

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 larger segments, and road parts are extracted from these segments. Roads in the image are often covered by several extracted 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 after gap weights are determined. A digital surface model, if available, is used in the grouping and road extraction step in order to prevent high regions from being extracted as roads. The results of the road extraction with and without the digital surface model are compared in order to show how the extraction is improved by the surface model. It also shows what can still be expected from the extraction if no digital surface model is available.

1. INTRODUCTION

Roads are a very important part of the infrastructure, especially in urban areas. Road data are used in many applications, for example car navigation systems. For these applications it is important that the road data are up-to-date and correct. As the road network is subject to change, especially in suburban areas, the road databases have to be updated frequently. This is often done manually with the help of aerial or satellite images. In order to reduce the costs and the time required for map updating, it is desirable to use automatic procedures for the extraction of roads from these images. Today, roads are to a large degree still extracted manually, especially in urban areas, because of the relatively high complexity of urban environments compared to open landscapes. For open landscapes, road extraction algorithms that are reasonably reliable already exist, e.g. (Zhang, 2004). This was confirmed by the EuroSDR test on road extraction (Mayer et al., 2006). In this test, several state-of-the-art methods for road extraction were compared, using imagery with a resolution of 0.5-1.0 m. The results were reasonably good in rural scenes of medium complexity, but the algorithms did not perform well in urban or suburban areas.

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. Road extraction algorithms can be classified into line-based approaches and region-based approaches. Line-based approaches, which model roads as one-dimensional linear objects, are mainly used in open landscapes with images of middle to low resolution, and they are not suitable for urban areas. An approach for urban areas that extracts middle lines and edges of roads and groups them to form road lanes using aerial images of very high resolution (0.1 m) is described by Hinz (2004). In most other approaches regions are extracted from images with a resolution of

approximately 1 m. One example is (Zhang and Couloigner, 2006), where a colour image is classified and the regions classified as roads are refined in order to separate roads from false positives such as parking lots. Another example for a region-based approach is (Hu et al., 2007), where footprints of roads are extracted based on shape, and the roads between the footprints are tracked. 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 (Zhang and Couloigner, 2006; Doucette et al. 2004), Digital Surface Models (DSMs) (Hinz, 2004) or both (Hu et al., 2004). Information about the position of roads from an existing road database can also be used, e.g. (Mena and Malpica, 2005). Prior information about the road network is another possible source of information. Price (1999) assumes that the road network forms a regular grid. This is also done by Youn and Bethel (2004), though they use less strict requirements for the grids.

In this paper, a region-based approach for road extraction from aerial colour images with a resolution of 0.1 m is presented. Optionally, a DSM can be used as an additional source of information. Apart from the DSM, our approach does not require other sources of information such as an existing database, as used in (Mena and Malpica, 2005). Since we work in suburban areas, the approach does not rely on particular properties of roads like road markings, as used in (Hinz, 2004) or a regular road grid, as used in (Price, 1999), and all roads should be extracted, not only major roads. In the approach, an image is first segmented and then road parts are extracted from the segments. These road parts are assembled into road subgraphs. In this way, there is no need to assume that a whole road can be extracted undisturbed. The subgraphs can contain different branches which represent different hypotheses for the

course of the roads. In order to find the most probable course of the road, the subgraphs are evaluated using relations between the road parts and linear programming. If a DSM is available, it can be used in the grouping and road part extraction processes. DSMs have been used in the past for road extraction, but their influence was limited due to the relatively poor performance of standard image matching techniques (Zhang, 2004). We think that with the advent of new dense image matching techniques, e.g. (Pierrot-Deseilligny and Paparoditis, 2006; Hirschmüller, 2008), the importance of incorporating 3D information into the road extraction process will increase. In this paper, the extraction results that were achieved with and without the DSM are compared in order to demonstrate the respective potentials for road extraction. The main goal of this paper is to present the new method for road subgraph evaluation and to assess the influence of the DSM on the road extraction results. The road extraction approach is described in Section 2. The segmentation and road part extraction, which are explained in detail in (Grote et al., 2007; Grote and Heipke, 2008), are only reviewed briefly. Our new method for road subgraph evaluation is discussed in more detail, as well as the incorporation of the DSM. In Section 3, results are presented with a comparison between the results achieved with and without the DSM. Section 4 gives conclusions and directions for future work.

2. APPROACH

2.1 Overview

Our goal is the extraction of roads from high resolution aerial images in suburban areas. We use colour infrared (CIR) images with a ground resolution of approximately 10 cm. Optionally, a DSM, e.g. generated by image matching, can also be used. The approach consists of three steps, namely segmentation, road part extraction and subgraph generation. In the segmentation step, the image is first divided into many small segments, which are then grouped into larger segments having meaningful shapes. Potential road parts are extracted from the grouped segments using shape criteria. If a DSM is available, height can be used as additional criterion in the grouping and road part extraction steps. The road parts are then assembled to road subgraphs (Fig. 1) if they potentially belong to the same road; junctions are not considered in this step. Several branches are allowed to be present in one subgraph. In the next step, these ambiguities are resolved by optimising the graph in a way that finds the best possibility for the course of the road without branches.




Figure 1. Road subgraphs. Dashed lines: real road network; grey rectangles: extracted road parts; continuous lines: edges of road subgraphs. The blue lines delineate two examples for distinct road subgraphs.

In Fig. 1, the term *road subgraph* and its components are explained. The term *subgraph* is used in order to indicate that it does not represent a complete, interconnected road network. A road subgraph consists of several assembled *road parts*. A road part is a segment which is classified as a road. It can correspond to a whole road between two junctions or only a part of the road, or it can be a false positive. Each subgraph extends only as far as road parts can be found in a more or less straight continuation; in this way, each subgraph usually represents only one road. Each road part in a subgraph has two *nodes* which are connected via a *road edge*. A node can also maintain connections to nodes of other road parts via *gap edges*. These gap edges can be understood as hypotheses for connections between extracted road parts that were missed in the original road part extraction process. If more than one such connection exists at one node, the node has several *branches*. These branches correspond to conflicting hypotheses for a completion of the road. In order to achieve a consistent road network, these conflicts have to be resolved by road subgraph evaluation.

2.2 Segmentation and Road Part Extraction

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 (Shi and Malik, 2000) 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. Details can be found in (Grote et al., 2007).

The normalized cuts algorithm results in a considerable oversegmentation. This is necessary in order to preserve most road borders, but as a result, the initial segments must be grouped in order to obtain segments that correspond to road parts. Grouping is carried out iteratively using colour and edge criteria, this time considering the properties of the regions (as opposed to those of the pixels, which were used in the initial segmentation). Segments with irregular shapes that cover roads across junctions can occur in this step. Therefore, the skeletons of the segments are examined. If they have several long branches (not to be confused with the branches of subgraphs), the segments are split.

In the next step, hypotheses for road parts are extracted from the grouped segments. Geometric and radiometric criteria are used for the extraction. The geometric criteria are elongation (ratio of squared perimeter to area), width constancy (ratio of mean width to standard deviation) and difference to average road width. As radiometric criteria, the NDVI (normalized difference vegetation index) and the standard deviation of colour are used. In addition, dark areas are excluded because shadow areas often have similar geometric properties to road parts. The parameters used for the experiments described in this paper are listed in Table 1. The elongation, width constancy, compliance with average road width and the NDVI are used to determine a quality measure for each road part hypothesis. The road parts are represented as regions; for the following road subgraph generation a representation by the centre lines and average widths is also used. For calculating the centre line, the region boundary is split into two parts at the points on the boundary that are farthest away from each other. Distance transforms are calculated for both parts, and the points where

both distance transforms within the road part have the same values make up the centre line. Further details of the road part extraction are explained in (Grote and Heipke, 2008).

min. elongation	70
width	3 m – 16 m
min. width constancy	1.7
min. intensity	40
max. NDVI	0
max std. deviation of colour	50

Table 1. Parameters for road part extraction.

2.3 Road Subgraph Generation and Evaluation

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 quality measure from the road part extraction. The criteria used to decide whether two road parts belong to the same road are the distance between the segments, the direction difference and the continuation smoothness. The reference points for the measurement of the direction difference and the continuation smoothness are the intersection points between the centre 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. However, 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 where the angles are measured.




Figure 2. Continuation smoothness.

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. The parameters used for the experiments described in this paper are shown in Table 2. 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. 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 quality measure among the remaining road parts until all road parts have been examined.

max. distance	50 m
max. direction difference	40°
max. smoothness angle	40°

Table 2. Parameters for road subgraph generation.

In most cases the 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. It is the goal of road subgraph evaluation to determine the best hypotheses, i.e. the hypotheses that are most likely to actually correspond to roads, and to discard all the other hypotheses. This goal is achieved 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 maximised or minimised (Dantzig, 1963). 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 optimised are restricted by hard constraints, which can be described by equations or by inequalities. In our case the constraints are inequalities resulting from the condition that no node of a subgraph should be connected to more than one gap edge after the optimisation. The objective function which is to be maximised is

$$w_1x_1 + \dots + w_nx_n \rightarrow \max \quad (1)$$

where $w_1 \dots w_n$ are weights of the gap edges that reflect the plausibility that the two road parts belong to the same road. The unknown variables are $x_1 \dots x_n$. There is one such unknown for each gap edge in the road subgraph. Each of these binary variables indicates whether its corresponding 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 maximisation under the constraints that each node i can only be associated to one gap edge:

$$\sum_{j \in E_i} x_j \leq 1. \quad (2)$$

E_i is the set of gap edges belonging to node i . The optimisation is carried out using the simplex method (Dantzig, 1963). 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. Section 2.2) 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. 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 normalised such that their sum equals 1.

After solving the linear programming problem, gap edges whose corresponding unknowns were determined to be 0 are discarded. This results in consistent road subgraphs that correspond to roads, which are considered to be the results of road extraction. However, all road parts are kept at this stage, so the falsely extracted road parts must be removed during the road network formation, which is still under development.

2.4 Including the Digital Surface Model

The road part extraction can produce false positives among segments having properties that are similar to road segments. False positives can disturb the later steps of linking the road parts and forming a road network. To avoid this, they have to be sorted out later when the network is formed or they have to be prevented from being extracted. The majority of the falsely extracted road parts are buildings, which can lead to subgraphs consisting only of false positives because buildings are often arranged in rows. As the most distinctive property of buildings that distinguishes them from roads is their height, a DSM can provide valuable additional information.

The DSM is employed in the grouping and road part extraction steps. It is not used for the initial segmentation which operates at pixel level because DSM inaccuracies in shadows and alignment errors caused by the fact that orthophotos are generated using a Digital Terrain Model (DTM) would affect the results adversely. In the grouping step the DSM is used to prevent segments with different heights from being merged. For this purpose, the differences of the mean heights are added to the grouping criteria. If the difference is larger than a threshold, the segments are not merged. The threshold is empirically determined; in our examples it is set to 1.5 m. This prevents building segments from being merged with road segments but allows for smaller height variations in the background.

For the road part extraction the DSM is used to prevent high objects from being extracted as roads. For this purpose, a normalised DSM (nDSM) representing objects above ground is determined. A coarse Digital Terrain Model (DTM) is generated from the DSM by morphological grey value opening. The DTM is then subtracted from the DSM, which yields the nDSM (Weidner and Förstner, 1995). The mean heights of the segments obtained from the nDSM are compared to a threshold. It was found that a threshold of about 1 m reliably distinguishes building parts and road parts. This threshold is used as additional criterion in the road part extraction.

3. RESULTS

The approach was tested on three subsets of an image showing a suburban scene from Grangemouth, Scotland. The image is a CIR orthoimage with a resolution of 10 cm. The data set also contains a DSM that was generated by image matching at a resolution of 20 cm in position and 10 cm in elevation. Elevated objects are represented well in the DSM, though unfortunately it is not known which method was used for its generation. For the three subsets, results of the road part extraction and road subgraph generation are presented, first obtained from the image data alone, and then from additionally using the DSM.

3.1 Results without DSM

Segmentation, grouping and road part extraction were carried out as described in Section 2.2 and (Grote et al., 2007; Grote and Heipke, 2008). Figures 3, 4 and 5 show the results of the road part extraction for the image subsets 1, 2 and 3, respectively. Whereas in subsets 1 and 2 most parts of the road network were extracted, significant parts of the road network are missed in subset 3. Each subset contains false positives, which are mainly found on buildings because buildings have similar radiometric and geometric properties to road parts. The results of the road part extraction were compared to manually extracted road regions. The manually extracted regions include areas occluded by shadows or trees, but exclude pavements. The completeness and correctness of the road parts computed according to (Heipke et al., 1997) are displayed in Table 3. They were determined on a per-pixel level and thus refer to the extracted areas. Table 3 shows that about two thirds of the road area could be detected, but almost half of the area classified as road area consists of false positives.

	Completeness	Correctness
Subset 1	66 %	57 %
Subset 2	89 %	59 %
Subset 3	31 %	49 %
Total	62 %	55 %

Table 3. Evaluation of road part extraction without a DSM.



Figure 3. Road parts extracted in subset 1 (yellow).



Figure 4. Road parts extracted in subset 2 (yellow).



Figure 5. Road parts extracted in subset 3 (yellow).

The road parts are then assembled into road subgraphs as is shown in Fig. 6 for subset 1. There are three subgraphs which contain different hypotheses; these are resolved using linear programming, as described in Section 2.3. In Fig. 7, only these three subgraphs are shown with the edges that are removed displayed in red. The results show that the optimisation favours connections between road parts that are similar in colour and width and maintain a more or less straight continuation.



Figure 6. Road subgraphs (without DSM) for subset 1. Different colours represent different road subgraphs.



Figure 7. Road subgraph evaluation (without DSM) for subset 1. Discarded gap edges are displayed in red.

3.2 Results with DSM

The grouping and the road part extraction were repeated using the DSM as additional information, as described in Section 2.4. Figures 8, 9 and 10 show the results of the road part extraction with the DSM for the image subsets 1, 2 and 3, respectively. Both the completeness and the correctness values (Table 4) have notably improved compared to the results without the DSM. The highest improvement in completeness is observed in subset 3; almost all roads are now covered with road parts for the greater part of their area. The highest improvement in correctness is observed in subset 1 where several buildings were extracted without the DSM.

	Completeness	Correctness
Subset 1	73 %	73 %
Subset 2	91 %	65 %
Subset 3	45 %	57 %
Total	70 %	65 %

Table 4. Evaluation of road part extraction with a DSM.

The subgraph generation and evaluation is conducted in the same way as before. The subgraphs for subset 1 can be seen in Fig. 11. Now there is only one subgraph with several branches, because the use of the DSM prevented some buildings from being extracted. The result of the evaluation of the remaining subgraph with branches is shown in Fig. 12.



Figure 8. Road parts extracted in subset 1 with DSM (yellow).



Figure 9. Road parts extracted in subset 2 with DSM (yellow).



Figure 10. Road parts extracted in subset 3 with DSM (yellow).

Compared to the visual impression of the extracted roads, the completeness and correctness values are relatively low. The computed correctness suffers from leakage at the borders of the road parts and from the fact that pavements, which are often extracted as roads, are not included in the reference data. The computed completeness would also be increased by constructing road parts corresponding to the gap edges that were accepted in subgraph evaluation.



Figure 11. Road subgraphs (with DSM), subset 1. Different colours represent different road subgraphs.



Figure 12. Road subgraph evaluation (with DSM), subset 1.

4. CONCLUSIONS

In this paper, an approach for the extraction of roads in suburban areas was presented, with the focus on a comparison between the extraction results achieved for image data alone and the results achieved for using a DSM as an additional information source. Our results show that the approach is suitable for the extraction of roads in suburban areas. The majority of roads can be detected even without a DSM, though there is a relatively high number of false positives, mostly buildings. Using a DSM improves both the completeness and the correctness of the results, primarily because buildings can now be clearly separated from roads. The correctness is improved because buildings are not extracted as false positives. The completeness is improved because incorporating the DSM into the grouping process provides a better grouping result from which more road parts can be extracted. Without a DSM, there are more subgraphs containing several branches, so that the importance of the subgraph evaluation is higher. The potential to find the real course of the road based on an optimisation of the interrelations between the road parts is shown in Figures 6 and 7. Subgraph evaluation can thus compensate for the lack of height information in the road part extraction stage. However, the improvement caused by using the height information in the grouping phase cannot be compensated. Road parts that remain undetected due to a poor performance of grouping based on image data alone cannot be detected at a later stage. Using a DSM thus certainly improves the quality of the results. This can be seen in particular for subset 3 (Fig. 5 vs. Fig. 10).

The road extraction process can still be improved in several ways. The parameters used for grouping, road part extraction and road subgraph generation could be learned from training samples, which probably would improve stability in different settings. The road extraction can also be improved by incorporating context objects such as trees, buildings and vehicles. It is planned to incorporate context objects into the evaluation of the gaps within the subgraphs, combined with the

interrelations described in this paper. Context objects can also be beneficial in the next steps, which include the formation of a road network by searching for junction hypotheses between road strings and removing isolated (mainly falsely extracted) road parts. The completeness and correctness values given in this paper were obtained from the single road parts. They are likely to improve during the network generation because the majority of falsely extracted road parts can be removed and the gaps between road parts in a string can also be counted as road.

ACKNOWLEDGEMENTS

This project is funded by the DFG (German Research Foundation) under grant HE 1822/21-1. 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 www.seas.upenn.edu/~timothee/software_ncut/software.html (accessed March 2009). The calculation of the linear program was made with the MILP solver `lp_solve` (lpsolve.sourceforge.net/5.5/, accessed March 2009).

REFERENCES

- Dantzig, G.B., 1963. Linear programming and extensions. Princeton University Press, Princeton, New Jersey, USA.
- Doucette, P., Agouris, P. and Stefanidis, A., 2004. Automated road extraction from high resolution multispectral imagery. *PE & RS* 70(12), pp. 1405-1416.
- Grote, A., Butenuth, M. and Heipke, C., 2007. Road extraction in suburban areas based on normalized cuts. In: *IAPRSIS XXXVI-3/W49A*, pp. 51-56.
- Grote, A. and Heipke, C., 2008. Road extraction for the update of road databases in suburban areas. In: *IAPRSIS XXXVII-B3b*, pp. 563-568.
- Heipke, C., Mayer, H., Wiedemann, C. and Jamet, O., 1997. Evaluation of automatic road extraction. In: *IAPRS XXXII-3/2W3*, pp. 47-56.
- Hinz, S., 2004. Automatic road extraction in urban scenes – and beyond. In: *IAPRSIS XXXV-B3*, pp. 349-355.
- Hirschmüller, H., 2008. Stereo processing by semi-global matching and mutual information. *IEEE TPAMI* 30(2), pp. 328-341.
- Hu, J., Razdan, A., Femiani, J.C., Cui, M. and Wonka, P., 2007. Road network extraction and intersection detection from aerial images by tracking road footprints. *IEEE TGRS* 45(12), pp. 4144-4157.
- Hu, X., Tao, C.V. and Hu, Y., 2004. Automatic road extraction from dense urban area by integrated processing of high resolution imagery and LIDAR data. In: *IAPRSIS XXXV-B3*, pp. 288-292.
- Mayer, H., Hinz, S., Bacher, U., and Baltsavias, E., 2006. A test of automatic road extraction approaches. In: *IAPRSIS XXXVI-3*, pp. 209-214.
- Mena, J.B. and Malpica, J.A., 2005. An automatic method for road extraction in rural and semi-urban areas starting from high-resolution satellite imagery. *Pattern Recognition Letters* 26(9), pp. 1201-1220.
- Pierrot-Deseilligny, M. and Paparoditis, N., 2006. A multiresolution and optimization-based image matching approach: an application to surface reconstruction from SPOT-HRS stereo imagery. In: *IAPRSIS XXXVI-1/W41*, pp. 73-77.
- Price, K., 1999. Road grid extraction and verification. In: *IAPRS XXXII-3/2W5*, pp. 101-106.
- Shi, J. and Malik, J., 2000. Normalized cuts and image segmentation. *IEEE TPAMI* 22(8), pp. 888-905.
- Weidner, U. and Förstner, W. 1995. Towards automatic building reconstruction from high resolution digital elevation models. *ISPRS J. Photogr. & Rem. Sens.* 50(4), pp. 38-49.
- Youn, J. and Bethel, J.S., 2004. Adaptive snakes for urban road extraction. In: *IAPRSIS XXXV-B3*, pp. 465-470.
- Zhang, C., 2004. Towards an operational system for automated updating of road databases by integration of imagery and geodata. *ISPRS J. Photogr. & Rem. Sens.* 58(3-4), pp. 166-186.
- Zhang, Q. and Couloigner, I., 2006. Automated road network extraction from high resolution multi-spectral imagery. In: *Proc. ASPRS Annual Conf.*, Reno, Nevada, 10 p., on CD-ROM.