

AUTOMATIC ROAD NETWORK EXTRACTION IN SUBURBAN AREAS FROM HIGH RESOLUTION AERIAL IMAGES

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: Road extraction, Urban, High resolution, Aerial, Automation

ABSTRACT:

In this paper a road network extraction algorithm for suburban areas is presented. The algorithm uses colour infrared (CIR) images and digital surface models (DSM). The CIR data allow a good separation between vegetation and roads. The image is first segmented in two steps: an initial segmentation using the normalized cuts algorithm and a subsequent grouping of the segments. Road parts are extracted from the segments and then first connected locally to form subgraphs, because roads are often not extracted as a whole due to disturbances in their appearance. Subgraphs can contain several branches, which are resolved by a subsequent optimisation. The optimisation uses criteria describing the relations between the road parts as well as context objects such as trees, vehicles and buildings. The resulting road strings, represented by their centre lines, are then connected to a road network by searching for junctions at the ends of the roads. Small isolated roads are eliminated because they are likely to be false extractions. Results are presented for three image subsets coming from two different data sets, and a quantitative analysis of the completeness and correctness is shown from nine image subsets from the two data sets. The results show that the approach is suitable for the extraction of roads in suburban areas from aerial images.

1. INTRODUCTION

As roads are an essential part of the infrastructure, accurate and up-to-date road databases are essential for many applications. Aerial and satellite images are often used to verify and update road databases manually, but it is desired to automate this process as far as possible in order to save time and costs. For open landscapes, there exist algorithms that work well enough for the practical application of database verification, e.g. Zhang (2004), Gerke & Heipke (2008), or update, e.g. Mena & Malpica (2005). In urban areas, the task is considerably more difficult because of the complex scene content.

Most automatic road extraction algorithms can be coarsely grouped into line-based approaches and region-based approaches. Line-based approaches are often used for road extraction in open landscapes, using images of relatively low resolution. An example for a line-based approach is presented by Baumgartner et al. (1999): lines are extracted in low resolution images. They are combined with edges extracted in high resolution images, and road segments are extracted where a line is bordered by two edges. The road segments are grouped iteratively into a road network. Similar approaches are also used in Wiedemann & Ebner (2000) and in Gerke & Heipke (2008). In urban areas, the usefulness of line-based extraction methods is limited because of the scene complexity. The scene contains many other objects with linear features. Line-based approaches for urban areas typically use some form of additional prior knowledge, constraints, or data sources. Frequently the road network in urban areas is assumed to be a fairly regular grid of straight roads. Long straight lines with little grey value variation (Shackelford & Davis, 2003) or with few crossing lines (Youn & Bethel, 2004) are searched for. Hu et al. (2004) use a digital surface model (DSM) from LIDAR as an additional data source to restrict the search space for the straight lines. These approaches usually cannot handle curved roads well.

In region-based approaches roads are modelled as elongated regions. Many approaches use multispectral classification as a first step to extract road areas or regions of interest for roads. For example, Zhang (2004) finds regions of interest for roads by a multispectral classification and by excluding high regions via a DSM; then parallel edges are extracted in the regions of interest. Roads are only extracted in the regions around database roads. Doucette et al. (2001) use hyperspectral data (210 channels) to extract road regions. Afterwards road pixels are grouped into a network with a k-median classification. Mena and Malpica (2005) use three different classification methods for colour and texture and combine them; the extracted regions are then vectorised. All these methods are developed for rural or at most semi-urban areas. There are only few region-based approaches for urban areas. Hinz and Baumgartner (2003), who work in dense urban areas, extract edges and ribbons that are assembled to road lane segments within regions of interest determined from a DSM. The lane segments are grouped into road segments and finally into a road network. Zhang and Couloigner (2006) first perform a multispectral classification and then filter the regions of the road class according to shape criteria. Poullis and You (2010) use pixel curve information together with the pixel colour for a classification into road and non-road pixels with the graph cut algorithm. Another interesting region-based approach is Hu et al. (2007): footprints are extracted based on their shape and then tracked. Junction footprints are distinguished from ordinary road footprints; in this way, a whole network can be extracted. Post-processing is necessary to remove false extractions.

Road extraction is often difficult if other objects such as buildings or trees (*context objects*) are close to the road, disrupting the appearance of the road or occluding it. Some approaches explicitly model context objects and include them in their extraction strategy. In rural areas, context objects are considered by Gerke and Heipke (2008), who use rows of trees as additional hints for roads, and by Baumgartner et al. (1999),

who model occlusions, shadows and disruptions of roadsides by driveways. There are only few examples for the explicit use of context objects in urban areas. Approaches incorporating height data use high regions implicitly as context objects to exclude them from the search space, e.g. Hu et al. (2004). Hinz and Baumgartner (2003) explicitly model cars and buildings and their relations to roads in order to assist the road extraction.

In this paper, we present a new approach for road network extraction in suburban areas. In this context ‘suburban’ means areas with relatively low buildings and not as densely built-up as inner city centres. We use high resolution colour infrared (CIR) images and a DSM, but no prior information from road databases in order to be able to deal with regions where this information is not available. Unlike other authors we do not rely on specific road patterns, straight roads or the existence of road markings because of the variations in these respects that occur especially in newly built-up areas. Our approach is region-based, but we start with a segmentation of the image, not with a multispectral classification. In this way, the method can be more easily transferred to different regions and sensors. Knowledge about the appearance of roads in aerial images is used already in the segmentation, which is based on normalized cuts (Shi & Malik, 2000). Context objects such as buildings, vegetation and vehicles are used in the road extraction process, because they can cause disturbances in the appearance of roads, for example by occlusions. The network is created by linking the extracted roads via junction connections; isolated short roads are eliminated in this process. In section 2 the road extraction method is described. In section 3 results are presented as well as an analysis of their completeness and correctness. Section 4 gives some conclusions and suggestions for further work.

2. METHODS

The goal of our approach is the extraction of a road network in suburban areas. We follow a region-based strategy on high resolution aerial images and use road-specific knowledge from the segmentation through the whole process to the network linking. A DSM is used as additional information. Road extraction starts with an initial segmentation of the image. Afterwards, the segments are grouped, and road parts are extracted. The road parts are connected locally, and ambiguous connections (links from one end of a road part to more than one other road part) are resolved through optimisation. Afterwards, the locally connected road parts (road strings) are linked to a network by setting up junction connections.

2.1 Image segmentation, grouping and road part extraction

2.1.1 Initial segmentation: For the initial segmentation the normalized cuts algorithm is used (Shi & Malik, 2000), in which an image is represented as a graph and segmented considering similarities between pairs of pixels. The segment borders are optimised globally such that the similarity of pixels between segments is minimal while the similarity of pixels inside the same segment is maximal. A weight matrix is set up; the weights represent the similarities between the pixel pairs. The Laplacian matrix is calculated from the weight matrix, and eigenvectors are calculated from the Laplacian. After a discretisation the eigenvectors define the segmentation of the image: each eigenvector represents a segment. As the dimension of both the weight matrix and the Laplacian is (number of pixels)², computing the eigenvectors is only computationally tractable if the weight matrix is sparse. Thus, non-zero weights are only assigned to pixel pairs in a local neighbourhood. It is

an advantage of the normalised cuts method that model knowledge can be integrated into the segmentation via the definition of the weight matrix. In our application, the weights are based on several similarity criteria specifically designed to separate road areas from non-road areas. One criterion is the colour similarity; another is the existence and strength of edges between the pixels. The similarity values are combined to one weight w_{ij}^0 for each pixel pair i and j . More details on the definition of these weights can be found in Grote et al. (2007). More recently we have integrated a new criterion based on the normalised difference vegetation index (NDVI) in order to distinguish between pixels with vegetation and pixels without vegetation. A threshold is applied to the NDVI and a new similarity weight $w_{ij} = w_{NDVI} \cdot w_{ij}^0$ is determined for pixel pairs not belonging to the same NDVI region, with $0 < w_{NDVI} \ll 1$, so that their weights will be lowered considerably. Using the NDVI improves the separation between roads and vegetation significantly.

2.1.2 Grouping: The result of the normalized cuts algorithm is over-segmented, which is necessary to ensure that most road borders in the image will be segment borders. But for road extraction, the segments first have to be grouped to larger segments. Two segments can be merged if they fulfil certain criteria, based on the appearance of roads. As the road surface is usually homogeneous at least in sections, the difference of the colour histograms and the edge strength along the shared border are used as a grouping criterion. Other criteria are the convexity of the merged regions, the shared border length (absolute and relative to the segment border lengths), and the mean height difference (from the DSM). The grouping is done iteratively. In each iteration cycle all segment pairs are evaluated according to the grouping criteria. In order to decide if two segments are candidates for merging, the values for the criteria are combined using fuzzy sets and a set of rules, ensuring that segments can be merged not only if all criteria are fulfilled but also if one or two are poor. For example, if at least two of the edge, colour and convexity criteria are very good and the third is still good, the criterion for the relative border length can be disregarded. All segment pairs that are candidates for merging are sorted by the sum of the normalised values for the criteria. In each iteration cycle, the best 10 % of segment pairs are merged.

2.1.3 Road part extraction: Road parts are extracted from the grouped segments according to geometric and radiometric criteria. Geometrically, road parts are elongated and in most cases convex regions with a limited range of widths, so the elongation (ratio of squared perimeter to area), the convexity (ratio of segment area to area of convex hull) and the width constancy (ratio of mean width to standard deviation of width) should be high and the average width should lie within the range of typical road widths. The centre line for each road part is determined by a distance transform. The average width is twice the average distance between the centre line and the segment borders. Additionally, the road parts should have a minimum length, and they should lie in low regions in the normalised DSM. As radiometric criteria a low NDVI and a low standard deviation of the intensity are required, and the intensity should neither be very low nor very high. All criteria must meet certain thresholds for the region to be extracted as a road part, but some criteria are balanced against each other: to allow the extraction of curved road parts, the convexity can be lower than the threshold if the elongation is high. After the extraction, adjacent road parts are merged if they have similar directions and if the merged region also fulfils the criteria for road parts. A quality measure is calculated for each road part

from the elongation, width, width constancy, length and NDVI. More details can be found in Grote and Heipke (2008).

2.2 Road subgraph generation and evaluation

2.2.1 Road subgraph generation: Roads can rarely be extracted completely as one single road part because of disturbances in the appearance of the roads in the images, e.g. occlusions. Therefore, road parts (represented by the extracted regions and their centre lines) are linked locally into subgraphs before the network generation starts. A subgraph corresponds to a set of connected road segments and may contain competing hypotheses. Two road parts can be connected if they fulfil certain criteria regarding distance, direction difference and continuation smoothness. The distance, in absolute terms as well as relative to the lengths of both road parts, should be small; the direction difference (difference of main directions of the road parts) should also be small. The continuation smoothness should be high, which means that the lateral offset between the ends of two linked road parts should be low. Starting with the road part with the best quality measure from the previous step, the road parts are linked iteratively according to the criteria described above. Connections from one end of a road part to more than one other road part are permitted in this stage to preserve all possible linking hypotheses, so that several branches in the subgraphs can occur. If no more road parts can be added to the subgraph according to the criteria, the procedure restarts with the remaining road parts, again beginning with the road part that has the best quality measure.

2.2.2 Road subgraph evaluation: If a subgraph has more than one branch, these branches are considered to be competing road hypotheses, all except one caused by falsely extracted road parts. Thus, it is necessary to determine the best connection hypotheses which result in a road string without branches. This is done in a separate step after the subgraph linking because it is not always possible to decide locally which one of several connections is the best: further connections of the involved road parts to other road parts should be taken into account.

The optimisation of the subgraphs is performed by formulating the task as a linear programming problem and solving it. In linear programming a linear function whose variables are subject to linear constraints is maximised or minimised (Dantzig, 1963). The linear function to be optimised for the road subgraph evaluation is

$$w_1 x_1 + \dots + w_n x_n \rightarrow \max \quad (1)$$

where n is the number of connection hypotheses in the subgraph. Each of the weights w_i describes the quality of the connection hypothesis i , and each of the unknown binary variables x_i describes whether the connection hypothesis i should be kept ($x_i = 1$) or discarded ($x_i = 0$). For each end of each road part, one constraint is formulated. The constraints are defined by the condition that an end of a road part may only be connected to one other road part: the sum of all x_i for connections to one end of a road part must not be higher than 1. The weight w_i for a connection is composed of two parts: the relation weight, which is determined by the geometric and radiometric relations between the two connected road parts, and the context weight, which is determined by context objects that can be found in the gap between the road parts. The relation weight is calculated as the product of weights $w_R(f_j)$ determined from six features f_j , thus $j \in \{1 \dots 6\}$. Three features are already used in the subgraph generation: continuation smoothness, distance and direction difference. The other features are the

qualities of the connected road parts, the colour difference and the width difference. The weight functions $w_R(f_j)$ linearly map the feature values into the interval $[0, 1]$ such that a high value indicates a good connection with respect to that feature. All individual weights are multiplied to obtain the total relation weight.

The context weight w_C is determined using the context objects vehicle, tree, vegetated area, building, shadow and asphalt region. They are extracted automatically, using relatively simple algorithms since they are only used to aid the road extraction (Grote et al., 2009). For the determination of the context weight, the connection hypothesis is transformed into a road part hypothesis in the gap with the average width of the connected road parts. The context weight is composed of the context relation weight w_{CR} and the context occlusion weight w_{CO} . The context relation weight w_{CR} indicates whether context objects support or contradict a road part hypothesis. 14 relation categories and their corresponding relation values v_k are defined based on the type of context object, the position and the orientation relative to the road part hypothesis. The values v_k depend on the relevance of the relation for support or contradiction of a road part hypothesis. They lie in the interval $[-0.5, 0.5]$, except for relations strongly contradicting the road part hypothesis, for which v_k is set to -10. The large negative number indicates that the road part hypothesis is highly improbable, despite any supporting context objects. For example, the relation category *building on road* strongly contradicts a road part hypothesis, and thus is assigned the relation value -10. The context relation weight is calculated from all appearances of relation categories in a gap:

$$w_{CR} = \sum_{k=1}^r \sum_{m=1}^a v_k / m \quad (2)$$

where r is the number of the relation categories found in the gap, and a is the number of appearances for relation category k . The context occlusion weight w_{CO} indicates whether the context objects occlude the road, causing the extraction to fail. It is measured by the percentage of the area of the road part hypothesis that is covered by the context objects vehicle, tree and shadow. The context weight w_C is the sum of the context relation weight w_{CR} and the context occlusion weight w_{CO} .

A combination of the context weight and the relation weight yields the total weight w_i for each connection. If the context weight is negative (indicating a strongly contradicting relation such as a building on the road hypothesis) the total weight is set to 0 regardless of the interrelation weight. If the interrelation weight is 0, the connection is only kept if the context weight is very good. The context weight is disregarded for small gaps because it is not reliable for short connection hypotheses. Otherwise, the context weight and the interrelation weight are combined by calculating the mean. After calculating the weights, the connections to be kept are determined by solving Eq. 1 for the x_i . Only the connections where $x_i = 1$ are kept. The evaluation reduces the subgraphs to road strings without branches. Each road string receives a quality measure that is determined as the mean of the quality measures of the individual road parts. For more details on the definition of the weights refer to Grote et al. (2009).

2.3 Network generation

The last stage of the road network extraction is the connection of the road strings to a road network. For this purpose, the road strings are first converted such that they are represented by their

centre lines and average widths. The centre line of a road string is composed from the centre lines of the individual road parts and the connecting lines between them, and then approximated by a polygon approximation. The polygon approximation starts with a straight line between the end points of the centre line and iteratively adds vertices if the distance of the approximation to the original is too high. The average width is calculated from the average widths of the individual road parts. A road region is calculated from the centre line and the average width. The new centre lines are used for the connection of the road network.

In the first step of the network generation, pairs of parallel roads that lie close together are searched for. From such a parallel pair, only the road that has the better quality measure is kept. Additionally, roads that are overlapped by other roads for a substantial part of their road region (defined by their centre lines and average widths as described above) are also deleted. In this way, some false extractions can be eliminated before the actual network linking. After that, junction connections are searched for among the remaining roads. At the end of each road a search region is defined as a semicircle whose centre is the end point of the road centre line and which points in the direction of an extension of the road. The radius of the search circle depends on the quality measure of the road: a road with a good quality measure has a large search radius. If another road is found inside the search region, a junction connection is created. Depending on whether both roads are collinear or not, the junction connection is created in different ways. If the roads are collinear, i.e. have a small direction difference, they are connected if the end point of the second road lies inside the search region of the first, and the junction connection is the connection of the two end points. If the roads are not collinear, the junction connection is constructed from the extension of the first road, and, if necessary, from the extension of the second road, i.e. the junction connection is either the connection between one road and the intersection point on the other road or it is the short polygon connecting the two end points via the intersection point. Additionally, intersections between two roads are searched for. Junction connections are verified before they are accepted: if several competing junction connections (e.g. two parallel but not collinear roads) exist at the end of a road, only one is kept. The junction connections are evaluated according to their length: shorter junction connections are considered more reliable.

After the creation and verification of junction connections, the road network consists of one or more connected components. A connected component consists of at least one road and possibly junction connections. Connected components are checked for significance: the total length of all roads must be more than the total length of all junction connections, and the total length of the connected component (roads and junction hypotheses) must exceed a minimum. An exception for the last condition is made if at least two open ends of the connected component lie near the image border; then it is possible that the connected component belongs to a road network beyond the image border.

3. EXPERIMENTS

The approach was tested on two different data sets. The first data set consists of an orthophoto generated from a scanned CIR aerial image of a suburban area in Grangemouth (Scotland) with a resolution of 10 cm, and a DSM obtained from image matching with manual post processing. The second data set consists of orthophotos generated from digital CIR aerial images with a resolution of 8 cm and a DSM from LIDAR data

from the DGPF (German Association for Photogrammetry and Remote Sensing) test site at Vaihingen (Germany) (Cramer and Haala, 2009). In our tests, we used six image subsets from the Grangemouth scene and three subsets from the Vaihingen scenes, each of them depicting suburban scenes. The roads extracted by our method were compared to a reference to assess the completeness and the correctness of the extraction results.

3.1 Results

Figures 1, 2 and 3 display the results from the subgraph generation for three of the subsets used for evaluation. Figures 1 and 2 show subsets from Grangemouth, whereas Figure 3 is taken from the Vaihingen data set. The subgraphs consist of the extracted road parts and the connecting lines found during the subgraph generation and evaluation. The results show the subgraphs after the evaluation, which means the road subgraphs consist of only one road string each. The subgraphs are depicted in different colours; roads that belong to the same subgraph have the same colour.



Figure 1. Accepted road subgraphs, subset 1 (Grangemouth).



Figure 2. Accepted subgraphs, subset 2 (Grangemouth).

Large parts of the road network could be extracted as road parts. Areas where the extraction fails typically lie at the image border (most notably in subset 2), at sharp turns or where the appearance of the roads is disturbed by trees and shadows. False extractions are rare, thanks to the DSM; most of them are driveways or parking lots. After the subgraph generation, most road parts that lie on the same road are connected. In subset 2, two road parts were first connected across two buildings, but the connection was eliminated after the context object evaluation (white dashed line in Figure 2). One connection in subset 2 was missed (white dotted line in Figure 2).

Figures 4, 5 and 6 show the road centre lines for the three subsets presented in Figures 1-3 after the network generation. The network of subset 1 is complete; the networks of subsets 2 and 3 show some gaps where junction connections were not constructed because the distance between the end points was too large. Subset 1 and 2 preserve an isolated false extraction each, because these lie close to the image border. These were labelled as less reliable and could as well be eliminated. The geometric accuracy of the road centre lines often decreases towards the ends of road parts. This is caused by irregularities in the shapes of the originally extracted road parts, which cannot always be overcome by the polygon approximation of the road strings. Junctions which were constructed by two extensions of road parts also often show geometric inaccuracies.




Figure 3. Accepted road subgraphs, subset 3 (Vaihingen).



Figure 4. Road network, subset 1 (Grangemouth).

3.2 Analysis of completeness and correctness

For a quantitative analysis of the completeness and correctness of the road network, the extracted road centre lines were compared to manually extracted centre lines. The completeness and correctness values are calculated as described in Wiedemann et al. (1998). Centre lines are compared in buffers of approximate road width, such that an extracted centre line that lies close to the roadside but inside the road is labelled as correct. Completeness and correctness values refer to the lengths of centre lines: only those parts that lie inside the buffers contribute to the correctness and completeness. Table 1 contains the completeness and correctness values of the three subsets shown in Figures 1-6. The most complete network is attained in subset 1, which on the other hand has the lowest correctness. This lack in correctness is for the most part due to the extraction of driveways which are not contained in the

reference. Another source for extractions labelled as incorrect is the lacking geometric accuracy; real false extractions are rare.

	Completeness	Correctness	Quality
subset 1	96%	88%	85%
subset 2	79%	91%	73%
subset 3	84%	95%	80%

Table 1. Completeness and correctness for displayed subsets.



Figure 5. Road network, subset 2 (Grangemouth).




Figure 6. Road network, subset 3 (Vaihingen).

Table 2 shows a summary of the quantitative analysis for all nine examined image subsets. The correctness is quite good. It is better than 90% in most of the subsets. The lowest correctness was 85% in a subset from the Grangemouth data set, where some parking areas were extracted as roads. The highest correctness achieved was 98%, also in a Grangemouth subset. The completeness is typically somewhat lower. The lower completeness from the Vaihingen data set is mainly caused by the fact that some roads were occluded by trees or at least affected by the shadows caused by the trees to an extent that no road parts could be extracted there. Whereas fragmented road parts can be connected via context objects by our method, the analysis of context objects (cf. Section 2.2.2) cannot detect new road parts. Therefore, the lowest completeness was 51%, in a subset from the Vaihingen data set. The highest completeness achieved was 97%, in a subset from the Grangemouth data set.

	Completeness	Correctness	Quality
Grangemouth	83%	92%	77%
Vaihingen	63%	94%	61%
Total	73%	93%	69%

Table 2. Completeness and correctness for all subsets.

4. CONCLUSIONS AND OUTLOOK

In this paper, an approach for the extraction of a road network in suburban areas was presented. CIR images and a DSM were used to first segment an image, then extract road parts and connect them to finally form a road network. The results presented in this paper show that the approach is suitable for the extraction of a road network in a suburban scene. From all examined subsets, three quarters of the road network could be extracted, and more than 90% of the extracted roads were correct. The approach was tested on two different data sets (Grangemouth and Vaihingen). Despite the fact that the two data sets had quite different sensor characteristics, we used identical parameters for our road extraction algorithm, with the exception of the NDVI threshold that had to be adapted manually. This suggests that the parameter set is quite robust; however, a further sensitivity analysis would be desirable. Whereas the total completeness was lower in the Vaihingen data set (mainly because the examined subsets there contained more roads covered by trees), the correctness was consistently good, which shows that the approach can be used for images from different sensors and different suburban areas. An important aspect to be improved is the geometric accuracy. This concerns several parts of the algorithm. The extraction of the centre lines from the irregularly shaped road parts could be improved by a previous orientation-dependent smoothing of the road parts. The junctions could be more explicitly modelled and their verification could be enhanced by using context objects in a similar way to that used for the subgraphs. When the network is extracted, the geometric positions of the roads could be improved using a snake-based algorithm. The completeness of the network could be improved by a search for gaps in the network and an evaluation of these gaps, e.g. by examining valleys in the DSM starting from road ends.

ACKNOWLEDGEMENTS

The Vaihingen data set was provided by the German Association for Photogrammetry and Remote Sensing (DGPF): <http://www.ifp.uni-stuttgart.de/dgpf/DKEP-Allg.html>. The normalized cuts calculation was adapted from a Matlab program from T. Cour, S. Yu and J. Shi: http://www.seas.upenn.edu/~timothee/software_ncut/software.html (accessed May 2010). The linear program was calculated with the LP solver `lp_solve`: <http://lpsolve.sourceforge.net/5.5/> (accessed May 2010).

REFERENCES

- Baumgartner, A., Steger, C., Mayer, H., Eckstein, W. and Ebner, H., 1999. Automatic road extraction based on multi-scale, grouping, and context. *Photogramm. Eng. Rem. S.* 65(7), pp. 777-785.
- Cramer, M. and Haala, N., 2009. DGPF Project: Evaluation of digital photogrammetric aerial based imaging systems – overview and results from the pilot centre. In: IAPRSSIS XXXVIII(1-4-7/W5), on CD-ROM.
- Dantzig, G.B., 1963. *Linear programming and extensions*. Princeton University Press, Princeton, New Jersey, USA.
- Doucette, P., Agouris, P., Stefanidis, A. and Musavi, M., 2001. Self-organized clustering for road extraction in classified imagery. *ISPRS J. Photogramm.* 55(5-6), pp. 347-358.
- Gerke, M. and Heipke, C., 2008. Image-based quality assessment of road databases. *Int. J. Geogr. Inf. Sci.*, 22(8), pp. 871-894.
- Grote, A., Butenuth, M. and Heipke, C., 2007. Road extraction in suburban areas based on normalized cuts. In: IAPRSSIS XXXVI(3/W49A) pp. 51-56.
- Grote, A. and Heipke, C., 2008. Road extraction for the update of road databases in suburban areas. In: IAPRSSIS XXXVII(B3b), pp. 563-568.
- Grote, A., Heipke, C., Rottensteiner, F. and Meyer, H., 2009. Road extraction in suburban areas by region-based road subgraph extraction and evaluation. In: *Urban Remote Sensing Joint Event, Shanghai, 20-22 May*, on CD.
- Hinz, S. and Baumgartner, A., 2003. Automatic extraction of urban road networks from multi-view aerial imagery. *ISPRS J. Photogramm.* 58(1-2), pp. 83-98.
- 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: IAPRSSIS XXXV(B3), pp. 288-292.
- 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 TGARS* 45(12), pp. 4144-4156.
- 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 Recogn. Lett.* 26(9), pp. 1201-1220.
- Poullis, C. and You, S., 2010. Delineation and geometric modeling of road networks. *ISPRS J. Photogramm.* 65(2), pp. 165-181.
- Shi, J. and Malik, J., 2000. Normalized cuts and image segmentation. *IEEE TPAMI* 22(8), pp. 888-905.
- Shackelford, A.K. and Davis, C.H., 2003. Fully automated road network extraction from high-resolution satellite multispectral imagery. In: *Proc. IGARSS, Toulouse, Vol. 1*, pp. 461-463.
- Wiedemann, C., Heipke, C., Mayer, H. and Jamet, O., 1998. Empirical evaluation of automatically extracted road axes. In: K.J. Bowyer and P.J. Phillips (eds.): *Empirical Evaluation Methods in Computer Vision*, IEEE Computer Society Press, Los Alamitos, California, pp. 172-187.
- Wiedemann, C. and Ebner, H., 2000. Automatic completion and evaluation of road networks. In: IAPRSSIS XXXIII(B3/2) pp. 979-986.
- Youn, J. and Bethel, J.S., 2004. Adaptive snakes for urban road extraction. In: IAPRSSIS 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. Photogramm.* 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.