

# AUTOMATIC ROAD EXTRACTION IN URBAN SCENES — AND BEYOND

Stefan Hinz

Remote Sensing Technology, TU München, Arcisstr. 21, 80 333 München – Stefan.Hinz@bv.tum.de

## Commission III, Working Group 3

**KEY WORDS:** Image Understanding, Road Extraction, Fusion, Self-Diagnosis, Aerial Imagery

## ABSTRACT

In this paper, we present work on automatic road extraction from high resolution aerial imagery taken over urban areas. In order to deal with the high complexity of this type of scenes, we integrate detailed knowledge about roads and their context using explicitly formulated scale-dependent models. The knowledge about how and when certain parts of the road and context model are optimally exploited is condensed in the extraction strategy. To exploit information from multiple views, a fusion strategy for road objects (e.g. lanes) has been developed. It is based on internally computed quality measures and embedded in the system's concept of self-diagnostic extraction algorithms. The analysis of the final results shows benefits but also remaining deficiencies of this approach. We give an outlook on the utilization of the approach in applications related with traffic monitoring in urban areas.

## 1 INTRODUCTION

From a practical point of view, research on automatic road extraction in urban areas is mainly motivated by the importance of geographic information systems (GIS) and the need for data acquisition and update for GIS. This demand is strikingly documented in the survey on 3D city models initiated by the European Organization for Experimental Photogrammetric Research (OEEPE, now called EuroSDR) a few years ago (Fuchs et al., 1998). Applications of road data of urban areas include analyses and simulations of traffic flow, estimation of air and noise pollution, street maintenance, etc.

From a scientific perspective, the extraction of roads in complex environments is one of the challenging issues in photogrammetry and computer vision, since many tasks related to automatic scene interpretation are involved. Factors greatly influencing the scene complexity are, for instance, the number of different objects, the amount of their interrelations, and the variability of both. Moreover, each factor—and thus the scene complexity—is related to a particular scale. To accommodate for such factors, techniques like detailed semantic modelling, contextual reasoning, and self-diagnosis have proven to be of great importance over the past years. It is clear that these techniques must be integral parts of an extraction system to attain reasonably good results over a variety of scenes.

Before describing the details of how our approach tries to cope with these challenges we briefly review work on automatic road extraction with emphasis on approaches dealing with urban environments (Section 2). In Section 3, we present underlying ideas and basic components of our road and context model. The extraction strategy is outlined in Section 4, illustrated by results of intermediate steps. Special focus is thereby on fusing road objects extracted in multiple overlapping images. In Section 5, a numerical evaluation of the results currently achievable with our system is given followed by a discussion of the advantages and remaining deficiencies of the proposed approach. We conclude the paper with an outlook on future work (Sect. 6).

## 2 RELATED WORK

Compared to the relatively high number of research groups focusing their work on road extraction in rural areas, only a few groups

work on the automatic extraction of roads in urban environments (see articles in (Grün et al., 1995, Grün et al., 1997, Agouris and Stefanidis, 1999, Baltsavias et al., 2001)). Most of the past and current efforts in road extraction rely on road models that describe the appearance of roads in rural terrain rather than in settlements. However, throughout all the different approaches, some issues have proved to be essential: By integrating a flexible, detailed road and context model one can capture the varying appearance of roads and the influence of background objects such as trees, buildings, and cars in complex scenes (Baumgartner et al., 1999, Ruskoné, 1996, Strat and Fischler, 1995). The fusion of different scales helps to eliminate isolated disturbances on the road while the fundamental structures are emphasized (Mayer and Steger, 1998). This can be supported by considering the function of roads connecting different sites and thereby forming a fairly dense and sometimes even regular network. Hence, exploiting the network characteristics adds global information and, thus, the selection of the correct hypotheses becomes easier (Fischler and Heller, 1998, Price, 2000, Wiedemann and Ebner, 2000). Last but not least, another important point is the integration of self-diagnosis techniques. They are used to evaluate the reliability of hypotheses of both low level features and higher level objects, which in turn facilitates decisions that inherently appear during the extraction process (Hinz and Baumgartner, 2002, Tupin et al., 1999, Tönjes et al., 1999).

## 3 MODEL

### 3.1 Road Model

The road model illustrated in Fig. 1 a) compiles knowledge about radiometric, geometric, and topological characteristics of urban roads in form of a hierarchical semantic net. The model represents the standard case, i.e., the appearance of roads is not affected by relations to other objects. It describes objects by means of "concepts", and is split into three levels defining different points of view. The real world level comprises the objects to be extracted: The road network, its junctions and road links, as well as their parts and specializations (road segments, lanes, markings,...). These concepts are connected to the concepts of the geometry and material level via concrete relations (Tönjes et al., 1999). The geometry and material level is an intermediate level which represents the 3D-shape of an object as well as its




Figure 1: Model for roads and their context.

material describing objects independently of sensor characteristics and viewpoint (Clément et al., 1993). In contrast, the image level, which is subdivided into coarse and fine scale, comprises the features to detect in the image: Lines, edges, homogeneous regions, etc. Whereas the fine scale gives detailed information, the coarse scale adds global information. Because of the abstraction in coarse scale, additional correct hypotheses for roads can be found and sometimes also false ones can be eliminated based on topological criteria, while details, like exact width and position of the lanes and markings, are integrated from fine scale. In this way the extraction benefits from both scales.

### 3.2 Context Model

The road model is extended by knowledge about context: So-called context objects, i.e., background objects like buildings or vehicles, may hinder road extraction if they are not modelled appropriately but they can substantially support the extraction if they are part of the road model. We define global and local context:

**Global context:** The motivation for employing global context stems from the observation that it is possible to classify semantically meaningful image regions—so-called context regions—where roads show typical prominent features and where certain relations between roads and background objects have a similar importance and characteristics. Consequently, the relevance of different components of the road model and the importance of different context relations (described below) must be adapted to the respective context region. In urban areas, for instance, relations between vehicles and roads are more important since traffic is usually much denser inside of settlements than in rural areas. We distinguish urban, forest, and rural context regions, which are extracted by a texture-based segmentation (see (Baumgartner et al., 1999, Hinz et al., 2001)).

**Local context:** We model the local context with so-called context relations, i.e., certain relations between a small number of road and context objects. In dense settlements, for instance, the footprints of buildings are almost parallel to roads and they give therefore strong hints for road sides. Vice-versa, buildings or other high objects potentially occlude larger parts of a road or cast shadows on it. A context relation “shadow”, for instance, can tell an extraction algorithm to choose modified parameter settings. Also vehicles occlude the pavement of a lane segment. Hence, vehicle outlines as, e.g., detected by the algorithm presented in (Schlosser et al., 2003) can be directly treated as parts




Figure 2: Model for grouped markings:

| Components used for <b>extraction</b>                                                                                                                                                    | Components used for <b>evaluation</b>                                                                                                                                                                               |
|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <ul style="list-style-type: none"> <li>▷ Orientation difference of pairs of markings and gaps: limited</li> <li>▷ Gap length: limited</li> <li>▷ Length of group: lower bound</li> </ul> | <ul style="list-style-type: none"> <li>▷ Lengths of markings and gaps: constant within group</li> <li>▷ Overall curvature <math>\tilde{s}</math>: low</li> <li>▷ Height variation of <math>s</math>: low</li> </ul> |

of a lane. In a very similar way, relations to sub-structures and the integration of GIS-axes—though not used here—can be modelled. Figure 1 b) summarizes the relations between road objects, context objects, and sub-structures by using the concepts “Lane segment” and “Junction” as the basic entities of a road network.

### 3.3 Model for Self-Diagnosis

In order to enable the computation of internal quality measures during extraction, the criteria (i.e. model components) defining an object are divided into two different types. Model components of the first type are used to extract an instance of an object and the components of the second type serve for evaluating its quality. For guaranteeing an unbiased evaluation, model components belonging to different types should be independent from each other. Figure 2 illustrates this for the case of grouped markings: Components listed in the left column are used to create a group of markings. The parameters listed in the right column are computed from each marking group and matched to predefined evaluation functions according to fuzzy-set theory. The fuzzy aggregation of all matches yields a confidence value indicating the reliability of the extracted object. A description of all involved evaluation models can be found in (Hinz and Baumgartner, 2002).

## 4 EXTRACTION AND FUSION OF ROAD OBJECTS

### 4.1 Extraction of Road Objects — Overview

The extraction strategy inheres knowledge about how and when certain parts of the road and context model are optimally exploited, thereby being the basic control mechanism of the extraction process. It is subdivided into three levels (see Fig. 3): Context-based data analysis comprises the segmentation of the scene into the urban, rural, and forest area and the analysis of context relations. While road extraction in forest areas seems hardly possible without using additional sensors, e.g., infrared or LIDAR sensors, the extraction in rural areas may be performed with the system of (Baumgartner et al., 1999). In urban areas, extraction of salient roads includes the detection of homogeneous ribbons in coarse scale, collinear grouping thin bright lines, i.e. road markings, and the construction of lane segments from groups of road markings, road sides, and detected vehicles. The lane segments are further grouped into lanes, road segments, and roads. During road network completion, finally, gaps in the extraction are iteratively closed by hypothesizing and verifying connections between previously extracted roads. Similar to (Wiedemann and Ebner, 2000), local as well as global criteria exploiting the network characteristics are used. Figures 4 and 5 illustrate intermediate steps of extraction and Figs. 6 and 7 show typical results. For details regarding the extraction we refer the reader to (Hinz et al., 2001, Hinz and Baumgartner, 2002). The system described there extracts roads from a single image and uses a DSM and views from other images to circumvent occlusions. In contrast, the new version extracts roads from all available images and fuses them in object space. The next section focuses on this particular issue.




Figure 3: Extraction Strategy.




Figure 4: Examples of intermediate steps during road extraction.

### 4.2 Fusion of Road Objects

To exploit information from multiple views, an appropriate fusion strategy has been developed, which is especially suitable for complex environments like urban areas. It can be characterized by following features: 1) It is based on objects, i.e., parts of the road network such as lane segments and road segments 2) It is carried out in object space 3) It is embedded in the system's concept of self-diagnostic extraction algorithms. From a methodological viewpoint, the novelty of this approach mainly relates to the incorporation and use of self-diagnosis algorithms for fusion. The first two points, however, accommodate the special properties of urban scenes and are thus of no minor importance. In the following comments on each point are given:

**Ad 1)** Fusion is based on objects because, as mentioned above, aerial images of urban areas show very high complexity. If fusion would be based on low level image primitives like raw gray values or edge structures, either an extremely accurate DSM must be given (effectively a 3D city model) or the fusion algorithm has to cope with many ambiguities and many conflicting hypotheses that occur when matching primitives over different images. Hence, our philosophy is to stay in 2D as long as possible and to extract objects of large extent and high semantics in each image separately. Matching such kind of objects over images is much easier and the requirements on a DSM can be relaxed significantly. In the case of our road extraction system, the objects which are subject for fusion are lane segments extracted in each available image. These are constructed in previous processing steps from groups of markings (i.e., thin bright lines) and (anti-)parallel road sides (i.e., grayvalue edges) while constraining them to enclose a homogeneous region or alternatively a vehicle (Hinz et al., 2001, Schlosser et al., 2003).

**Ad 2)** The main reason for performing fusion in object space is its natural way in treating each image with equal importance and not preferring any image a priori. Thus a dependence of the final results on the processing order of the images can be avoided. As a side effect, objects extracted in images of different resolutions may be combined easily and all necessary parameters can be passed in real-world values.

**Ad 3)** The fusion algorithm is embedded in the system's concept of self-diagnostic extraction algorithms. The idea behind this approach is that each module used during extraction should attach its result with a confidence value indicating the quality how well the job has been done. Our approach to define evaluation criteria from which the confidence values can be calculated is to split up the components of the underlying object models into two different types. Model components of the first type are used for extracting an instance of an object and the components of the second type serve as criteria for evaluating the quality of the extracted instance. For guaranteeing an unbiased evaluation, model components belonging to different types should be independent from each other. In order to evaluate a certain object, pre-defined fuzzy functions are used. Since the road model underlying our




Figure 5: Lanes extracted in different images (a,b) and their superimposition (c) calculated using DSM (d: grayvalue coded DSM).

system is designed in a hierarchical way (roads consist of lanes which again consist of markings and road sides, etc.), the confidence measure of each object are used in three different ways:

a) *Confidence propagation*: Confidence values of lower level objects (e.g., groups of markings and road sides) are combined using the principles of fuzzy-set theory and propagated to the next level of the model hierarchy (e.g. lanes).

b) *Autonomous evaluation*: According to our model for self-diagnosis, at each level, object knowledge not used for extraction or evaluation at lower levels is incorporated, e.g., each lane should have a parallel counterpart (one lane roads are not considered). Note that this evaluation is independent of propagated confidence values (therefore "autonomous").

c) *Consistency check*: The score of autonomous evaluation of a higher level object are used to test the consistency of lower level objects. Consider, for instance, a hypothesis of a two-lane road segment (i.e., the higher level object) of which the first lane is extracted correctly but the other one is extracted only in fragments, e.g., due to inhomogeneities of the pavement. The latter lane hypothesis has consequently a low rating through autonomous evaluation, however, from the higher level point of view, there is strong evidence that this particular hypothesis is correct. Hence, such a hypothesis would pass the consistency check and is kept for further processing. In general, this means for the implementation that a hypothesis—regardless of its autonomous evaluation—is kept as long as the next level in the model hierarchy is completely processed and evaluated.

**Implementation:** This concept is also applied and implemented for fusion of road information from multiple views. Lanes are extracted in each image separately (see Fig. 5) and projected on a fairly accurate DSM (grid size and accuracy ca. 2m). In case of overlapping lanes, the lane having the best (propagated) confidence value is selected first and its mutual overlap with other lanes is computed. The score for autonomous evaluation of such a lane is calculated from the overlap ratios of lanes extracted in other images including weights for their deviation in position and direction. After deleting redundant parts of lanes the lane with the second highest confidence value is selected, and so forth. Thus a unique set of fused lanes is achieved. In the next hierarchy level, road segments are constructed from the fused lanes, i.e., parallel and collinear lanes are merged. Note, that the individual lanes of a road segment may be fragmented as long as a parallel lane provides a connection from one lane fragment to another fragment. The average degree of fragmentation of a road segment serves as consistency check for the fused lanes, i.e., lanes are rejected if not enough evidence is given for grouping them into larger road objects.

Tests with less accurate DSMs have shown that the use of lanes as objects to be fused may lead to matching ambiguities. Hence, an alternative version of the system (Hinz, 2003) uses the object "road segment"—an object with more semantics (see the model hierarchy in Fig. 1 a)—for fusion.

## 5 EVALUATION OF THE RESULTS AND DISCUSSION


Figures 6 and 7 illustrate the final result of road extraction in two parts of the Zurich Hoengg dataset (Baltasvias et al., 2001). As can be seen, major parts of the road networks have been extracted in spite of the high complexity of the scenes. The system is able to detect shadowed road sections or road sections with rather dense traffic (see e.g. Fig. 7 a and b). The results have been evaluated by matching the extracted road axes to manually plotted reference data. Table 1 summarizes the numerical values according to the definition of (Wiedemann, 2003). As can be seen, we achieve a completeness of more than 75 % and a correctness of about 95 % regarding the extracted road axes that could be linked into a network. Also the evaluation of the network characteristics yields satisfying results since for all evaluation criteria (detour/shortcut factor, topological completeness, topological correctness) values close to the optimum are reached.

| Evaluation criteria          | Data set I: | Data set II: |
|------------------------------|-------------|--------------|
| Completeness [%]             | 76.6        | 81.6         |
| Correctness [%]              | 98.8        | 95.0         |
| RMS-Error [m]                | 1.3         | 2.5          |
| Mean detour factor [ ]       | 1.04        | 1.05         |
| Mean shortcut factor [ ]     | 0.95        | 0.95         |
| Topological completeness [%] | 100.0       | 84.0         |
| Topological correctness [%]  | 96.2        | 100.0        |


Table 1: External Evaluation of extracted road axes.

However, it must be noted that some of the lane segments have been missed or have been linked incorrectly (Fig. 7 b). This is most evident at complex road junctions and crossings in both image parts, where only spurious features for the construction of lanes have been extracted. Another obvious failure can be seen at the right branch of the junction in the central part of Data Set II (Fig. 7 a). The tram and trucks in the center of the road have been missed since our vehicle detection module is only able to extract vehicles similar to passenger cars. Thus, this particular road axis has been shifted to the lower part of the road where the implemented parts of the model fit much better. As a consequence, the RMS-value drops down from acceptable 1.3m in Data Set I to poor 2.5m in Data Set II. The interested reader may be referred to the much more exhaustive evaluation carried out in (Hinz, 2003).

In summary, the results indicate that the presented system extracts roads even in complex environments. An obvious deficiency exists in form of the missing detection capability for vehicle types as busses and trucks. However, the main bottleneck of our system is the (still) weak model for complex junctions. Hence, besides the aforementioned improvement of verifying connection hypotheses, one of our next steps will be directed towards the modelling and reliable detection of road junctions. As a final remark regarding the percentages of correctness and completeness we would like to mention that, in spite of the definitely encouraging results, it would be unfair to disregard the fact that these percentages can be achieved only due to the expertise of the system developers in setting the parameters correctly (as it is surely



(a) External evaluation of road axes.



(b) External evaluation of lanes.

Figure 6: External Evaluation of Data Set I: Reference matching extraction (bold); missed reference (thin).


true for every experimental fully-automatic system at present). In this field, we are still at the stage of fundamental research and there are still many questions left open and still many steps to go so that a state of maturity is reached to envisage a transition to operational use.

## 6 OUTLOOK — BEYOND ROAD EXTRACTION

In the last section of this paper, we will show that results like those obtained above can give valuable support for other applications. We exemplify this by two complementary approaches for monitoring traffic in urban areas. The first approach uses optical data similar to that used for road extraction, while the second one is designed to extract vehicles from thermal infrared data. In contrast to most related work on car detection, both approaches rely upon local as well as global features of vehicles.

### 6.1 Car Detection in Optical Imagery

To model a vehicle for high resolution optical data, a 3D-wireframe representation is used that describes the prominent geometric and radiometric features of cars including their shadow region. The radiometric part of the model is adaptive because, during extraction, the expected saliencies of various edge features



(a) Data Set II: Reference matching extraction (bold); missed reference (thin).



(b) Detailed views of extracted lanes.

Figure 7: Extraction and evaluation of Data Set II.

are automatically adjusted depending on viewing angle, vehicle color, and current illumination direction. The extraction is carried out by matching this model "top-down" to the image and evaluating the support found in the image. On global level, the detailed local description is extended by more generic knowledge about vehicles as they are often part of vehicle queues. Such groupings of vehicles are modelled by ribbons that exhibit the typical symmetries and spacings of vehicles over a larger distance. To make use of the supplementary properties of local as well as global features, the algorithms for vehicle detection and vehicle queue detection are run independently first. Then, the results of both are fused and queues with enough support from the detailed vehicle detection are selected and analyzed for rectangular blobs to recover vehicles missed during the previous steps (see Fig. 8 a). Details regarding the implementation of this approach can be found in (Hinz, 2004b).

Typical problems are posed by cars that are not part of a queue and whose sub-structures (hood, windshield, etc.) give not enough evidence for a successful detection. However, the integration of intermediate or final results of road extraction helps especially to find such cars, since the road information around a car now supplements the (missing) evidence of a car's sub-structures. Figures 8 b) and c) show an example of extracting a car between the ends of two lane segments.

### 6.2 Car Detection in Thermal Imagery


Compared to optical data, thermal imagery has generally a lower resolution and usually a worse noise level because of the higher sensitivity of the scanner. However, thermal sensors show also a number of advantages—most notably their night imaging capability and their potential to derive temperature and temperature differences of objects, thus allowing for inferences about the current activity of objects even if they are not moving. For these reasons, thermal imagery has become a very attractive alternative for monitoring vehicle activity.



(a) Extracted queue of cars (without road information).



(b) Connection (black) between extracted lanes (white) as region of interest.



(c) Extracted car (incl. substructures) within region of interest.

Figure 8: Extraction of cars in optical imagery.

Since vehicles appear only as small, elliptical, dark or bright "blobs" in these images (resolution about 1m), many other objects in urban areas exhibit a very similar appearance. Thus, a reasonable good system for vehicle detection from thermal imagery must make use of additional information. Knowledge about the appearance of cars as repetitive patterns in dense traffic situations or in filled parking lots provides such additional information. This kind of knowledge is used in the example shown in Fig. 9 (Details regarding the methodology of the approach can be found in (Hinz, 2004a)).

However, as can be also seen from Fig. 9, each car that is too far away from another car has been rejected during generating vehicle queues from the individual car hypotheses. To detect also isolated cars with high confidence, information about the road — or even better: the lane a car is driving on — needs to be included. Unfortunately, this information is almost impossible to extract from thermal imagery itself, since road sides are rarely visible therein (see Fig. 9). But clearly, road data from an external source are appropriate means to deliver the necessary information, and the road extraction system described above is able to serve as such a source. What is more, since it extracts not only road axes but also lanes, even the current state of the system will potentially be better suited for this particular application than common map data. Further developments and tests will show whether this expectation will be met.

## REFERENCES

Agouris, P. and Stefanidis, A. (eds), 1999. *Integrated Spatial Databases - Digital Images and GIS*. Springer-Verlag, Berlin.

Baltsavias, E., Grün, A. and van Gool, L. (eds), 2001. *Automatic Extraction of Man-Made Objects from Aerial and Space Images (III)*. Balkema Publishers, Lisse, The Netherlands.

Baumgartner, A., Steger, C., Mayer, H., Eckstein, W. and Ebner, H., 1999. Automatic Road Extraction Based on Multi-Scale, Grouping, and Context. *Photogrammetric Engineering and Remote Sensing* 65(7), pp. 777–785.

Clément, V., Giraudon, G., Houzelle, S. and Sandakly, F., 1993. Interpretation of Remotely Sensed Images in a Context of Multisensor Fusion Using a Multispecialist Architecture. *IEEE Transactions on Geoscience and Remote Sensing* 31(4), pp. 779–791.

Fischler, M. and Heller, A., 1998. Automated Techniques for Road Network Modeling. In: *Image Understanding Workshop '98*, Morgan Kaufmann Publishers, San Francisco, CA, pp. 501–516.




Figure 9: Extraction of cars in thermal imagery: Detected car hypotheses/"blobs" (white), hypotheses linked into queues (black).

Fuchs, C., Gülch, E. and Förstner, W., 1998. *OEEPE Survey on 3D-City Models*. OEEPE Official Publication 35, pp. 9–123.

Grün, A., Baltsavias, E. and Henricsson, O. (eds), 1997. *Automatic Extraction of Man-Made Objects from Aerial and Space Images (II)*. Birkhäuser Verlag, Basel.

Grün, A., Kuebler, O. and Agouris, P. (eds), 1995. *Automatic Extraction of Man-Made Objects from Aerial and Space Images*. Birkhäuser Verlag, Basel.

Hinz, S., 2003. *Automatische Extraktion urbaner Straßennetze aus Luftbildern*. Dissertation, Technische Universität München.

Hinz, S., 2004a. *Car Detection in Aerial Thermal Imagery by Local and Global Evidence Accumulation*. Technischer Bericht, PF-2004-01, Lehrstuhl für Methodik der Fernerkundung, Technische Universität München.

Hinz, S., 2004b. *Detection of Vehicles and Vehicle Queues in High Resolution Aerial Images*. *Photogrammetrie — Fernerkundung — Geoinformation (PFG)* 3, pp. 203–215.

Hinz, S. and Baumgartner, A., 2002. *Urban Road Net Extraction Integrating Internal Evaluation Models*. In: *International Archives of Photogrammetry, Remote Sensing and Spatial Information Sciences*, Vol. 34(3A), pp. 163–168.

Hinz, S., Baumgartner, A., Mayer, H., Wiedemann, C. and Ebner, H., 2001. Road extraction focussing on urban areas. In: (Baltsavias et al., 2001), pp. 255–265.

Mayer, H. and Steger, C., 1998. Scale-Space Events and Their Link to Abstraction for Road Extraction. *ISPRS Journal of Photogrammetry and Remote Sensing* 53(2), pp. 62–75.

Price, K., 2000. *Urban Street Grid Description and Verification*. In: *5th IEEE Workshop on Applications of Computer Vision*, pp. 148–154.

Ruskoné, R., 1996. *Road Network Automatic Extraction by Local Context Interpretation: application to the production of cartographic data*. PhD thesis, Université Marne-La-Vallée.

Schlosser, C., Reitberger, J. and Hinz, S., 2003. Automatic Car Detection in High Resolution Urban Scenes Based on an Adaptive 3D Model. In: *IEEE/ISPRS joint Workshop on Remote Sensing and Data Fusion over Urban Areas*.

Strat, T. and Fischler, M., 1995. The Role of Context in Computer Vision. In: *IEEE Workshop on Context-based Vision*, pp. 2–12.

Tönjes, R., Growe, S., Bückner, J. and Liedke, C.-E., 1999. Knowledge-Based Interpretation of Remote Sensing Images Using Semantic Nets. *Photogrammetric Engineering and Remote Sensing* 65(7), pp. 811–821.

Tupin, F., Bloch, I. and Maitre, H., 1999. A First Step Toward Automatic Interpretation of SAR Images Using Evidential Fusion of Several Structure Detectors. *IEEE Transactions on Geoscience and Remote Sensing* 37(3), pp. 1327–1343.

Wiedemann, C., 2003. *External Evaluation of Road Networks*. In: *International Archives of Photogrammetry, Remote Sensing and Spatial Information Sciences*, Vol. 34(3W/8), pp. 93–98.

Wiedemann, C. and Ebner, H., 2000. *Automatic Completion and Evaluation of Road Networks*. In: *International Archives of Photogrammetry and Remote Sensing*, Vol. 33(B3/2), pp. 979–986.