

ROAD EXTRACTION IN URBAN AREAS SUPPORTED BY CONTEXT OBJECTS

Stefan Hinz, Albert Baumgartner

Technische Universität München, D-80290 Munich, Germany

Tel: +49-89-2892-3880

Fax: +49-89-280-9573

E-mail: {hinz | albert}@photo.verm.tu-muenchen.de

KEY WORDS: Image Understanding, Road Extraction, Context

ABSTRACT

In this contribution, we introduce our concept for automatic road extraction in urban areas using high resolution aerial imagery and dense height information. For facilitating the extraction we use a hierarchical road model comprising different road elements (markings, lanes, junctions, road network) at different scales as well as context objects of roads like buildings and vehicles. Global contextual knowledge about roads in residential areas helps us to focus on certain image parts, and thus, to cope with the inherently high complexity of urban scenes. Road extraction is then performed in a mainly data driven fashion starting with the extraction of the basic road elements at the lowest level of the road model and reaching the top level, i.e., a topologically consistent road network, with the final steps of processing. On each step in between, hypotheses for the respective road elements are generated and validated, and thus, evidence about roads is collected and transferred to the next step. Since urban roads are often characterized by rather dense traffic the detection of vehicles and vehicle convoys plays a major role. The results of the currently implemented modules give us rise to further realize this concept.

1 INTRODUCTION

In the past, the automation of road extraction from digital imagery has received considerable attention. Research on this issue is mainly motivated by the increasing importance of geographic information systems (GIS) and the need for data acquisition and update for GIS. Especially in the field of urban planning, there is a high demand for actual, accurate and detailed information about the road network as required for applications like traffic flow analysis and simulation, estimation of air and noise pollution, street maintenance, etc.. The OEEPE Survey on 3D City Models by the European Organization for Experimental Photogrammetric Research (OEEPE) confirmed this demand yielding about 85% of the participants mentioning that information about the road network is of their greatest interest (see (Fuchs et al., 1998)).

For a number of cities in Europe and Northern America, larger parts of the road network are already digitally mapped, most important in commercial and military navigation and route planning systems. In practice, however, the level of detail of the data acquisition, i.e., which elements of a road are thought to be important and which are negligible, is (naturally) influenced by the application for which the data is collected. Since, for navigation purposes, the topology of the road network is on principle more important than the number of lanes (except traffic-dependent route planning), most existing area-wide road databases provide road information in form of the road axis attached with the road class or the road width. They thus lack of the proper level of detail mandatory for the above-mentioned applications (though, in many cases, the underlying data models would allow for a more detailed road description). Since a high percentage of the relevant urban features, including roads and their sub-structures, i.e., lanes and road markings, can be extracted from aerial images with high accuracy (Englisch and Heipke, 1998), the interpretation of aerial images is despite of existing GIS data a crucial task for the acquisition of detailed road data. E.g., (Bogenberger et al., 1999) use aerial images in order to count lanes and vehicles, measure queue lengths, and estimate vehicle velocities for calculating and simulating traffic flow conditions on high capacity roads. The time- and cost intensive manual procedure necessary for extracting the desired information constitutes the main bottleneck, which needs to be overcome.

In this paper, we present our concept for the automatic extraction of roads and their sub-structures from aerial imagery taken over urban areas. Besides multiple overlapping images, the approach relies on accurate height information as, e.g., derived from airborne laser data. The approach is designed to include additional information sources, without depending on it. Such sources are most important color channels, external information provided by GIS and results of semi- or fully-automatic building extraction. As in our previous work (Baumgartner et al., 1999, Baumgartner et al., 1997), we distinguish three "global contexts" called *context regions*: rural, forest, and urban. For road extraction in urban areas we use a model that describes roads and their sub-structures at different scales for different sensors. The extraction starts with the segmentation of Regions of Interest (RoI) based on height and image data at coarse scale. Then, the basic road

features at the lowest level of the road model, i.e., road markings and lanes are extracted separately in each fine scale image. After a first validation, which includes also the detection of vehicle, the results of the different images are fused into larger road segments (possibly consisting of several lanes). During the next step junctions and connection hypotheses between the road segments are generated and validated in each image. The results are fused again, and the next iteration starts. The extraction is completed if no more connections can be formed or verified.

The next section describes existing approaches on road extraction being most relevant to our work. After a short discussion of the appearance of roads in urban areas, our road model is introduced in Sect. 3. In Sect. 4 the extraction strategy is outlined in detail and results of the currently implemented modules are shown. Finally, the results are discussed and conclusions for our future work are drawn (Sect. 5).

2 RELATED WORK

Compared to the intense research activities on the extraction and visualization of buildings, site models, or city models (see (Mayer, 1999) and (Förstner, 1999) for an overview), only few groups work on the automatic extraction of roads in urban environments. One reason for this is that most of the past and actual efforts in road extraction, including our own previous work, rely on road models that describe the appearance of roads in rural terrain. Depending on the image resolution, roads are modeled as line-like structures (Wiedemann and Hinz, 1999, Heller et al., 1998, Gruen and Li, 1997, Geman and Jedynak, 1996) or relatively homogeneous areas satisfying certain shape and size features (Harvey, 1999, Zhang and Baltsavias, 1999, Baumgartner et al., 1999, Mayer et al., 1998, Trinder and Wang, 1998). Throughout all the different approaches some issues have proved to be of great importance: The fusion of different scales helps to eliminate isolated disturbances on the road while the fundamental structures are emphasized (Mayer and Steger, 1998). Exploiting the network characteristics adds global information and, thus, the selection of the correct hypotheses becomes easier (Wiedemann, 1999, Heller et al., 1998). The integration of context helps to cope with the influence of background objects like trees and buildings on the appearance of roads. In the following, we discuss three selected approaches, each consisting of certain promising elements, and each having significant influence on our concept:

(Vosselman and de Gunst, 1997) and (de Gunst, 1996) use a detailed and scale dependent road model for the actualization of outdated road maps. According to road construction guidelines, they model roads and freeways as aggregation of different lanes. Bright markings separate the individual lanes. Intersecting roads form junctions of different types (crossings, Y-junctions, and fly-overs). In order to detect changes in the road network, in a first step, the old information is validated using road features extracted from medium or small scale aerial images (about 0.4 m / 1.6 m resolution). Then, in case of inconsistencies, a change in the road width, e.g. an additional lane, or a junction where a new road branches off is hypothesized. Again, extracted road features verify the hypotheses. With this strategy it is possible to detect both changes in existing roads and completely new roads. They use, however, rather simple thresholding and profile matching techniques for the extraction of road features. The approach is consequently very sensitive for disturbances like cars, shadows or occlusions. Hence, this system seems only applicable to roads and freeways in open and rural areas, though the concept is more generic.

In the approach of (Ruskoné, 1996, Ruskoné et al., 1994, Airault et al., 1994), the interpretation of local context is used to verify an automatically extracted road network. The extraction starts with detecting seed points in a medium (about 0.5 m resolution) scale image followed by low level road tracking along homogeneous elongated areas. Then, hypotheses for the connection of the extracted road parts are generated and checked based on geometric criteria like distance and direction. The resulting road network is geometrically adjusted using snakes. For verification, the network is split into smaller pieces. A supervised classification assigns the meanings "road", "crossing", "shadow", "tree", or "field" to each piece. These so-called local contexts are validated using several geometric and radiometric criteria. In urban areas, where techniques based on road profiles or road sides may yield erratic results, the validation is done by detecting cars. With a neural network classifier car-like patterns are extracted and thereafter grouped into convoys (Ruskoné et al., 1996).

The approach of (Price, 1999) is particularly designed for extraction of urban street grids from medium and high resolution images (about 0.8 m - 0.2 m resolution). The road network is modeled as a combination of grids with a rather regular mesh size, the size of a single building block. Junctions – the nodes on the grid – are connected by individual road segments of approximately constant width and height. A human operator initializes the grid spacing and orientation by manual selection of three points specifying the first mesh. Then, the grid is iteratively expanded by adding new meshes. In each iteration, the new road segments are refined and evaluated by simultaneously matching their sides to image edges. Thus, longer portions of the road sides must be visible at least in one of different overlapping images. During final verification, height information and contextual knowledge are used to adjust the position of several consecutive road segments and to remove short road portions. To do this, the adjusted segments are evaluated anew but now regarding each segment's direct neighbors.

The above approaches show individually promising parts of road model and extraction strategy for different types of scenes. The varying appearance of roads in complex scenes can be captured by a detailed model for a road and its

contextual relations to background objects. Vice-versa, context helps to guide the interpretation, e.g. "easiest first", making the extraction easier to handle and more robust. This can be supported by considering the function of roads connecting different sites and thus forming a fairly dense and sometimes even regular network. Data from different and complementary sources, especially overlapping images and accurate height information, is very useful to address occlusions and shadows in urban environments. Furthermore, all of these knowledge sources need to be integrated into a single system.

3 MODELING URBAN ROADS

3.1 Appearance of Roads

Figure 1 visualizes two examples of urban roads. It is obvious that these images exhibit a more complex content than scenes showing rural areas since the number of different objects and their heterogeneity is much bigger. Generally, this implies that more details of the road and context model must be exploited for road extraction. In dense urban areas, for instance, some of the roads comprise several lanes that are linked by complex road crossings. What is more, by the increase of the number of objects the complexity of their relations grows, too. In Fig. 1a), for example, some parts of the streets are occluded by vehicles, especially at the road sides. Hence in this particular case, a road is mainly defined by groups of (parking) cars but not by parallel road sides or homogeneous surface. A similar relation is the occurrence of shadows cast by high buildings. A road generally appears bright in open areas, but in the case of shadows two problems for the extraction arise: (1) the surface is darkened significantly, and (2) at the margin of the shadow regions, strong gray value edges in almost any direction may occur on the road disturbing the usually homogeneous reflectance.

Figure 1b) shows a different kind of problem: The roof of the rectangular building in the center of the image could be wrongly identified as a parking lot because its shape and reflectance properties match the ones of a road-like object almost perfectly. Only the combination with height data as given by a DSM (Digital Surface Model) or, as in this case, implicitly given by a corresponding shadow region provides enough information for avoiding this misdetection.

What follows is, that on one hand those features of a road ought to be selected on which the influence of the above mentioned phenomena is minimal. On the other hand, it is very important to consider context objects, in particular different kinds of vehicles, in order to explain abnormal changes in the appearance of a road. Hence, the model described below consists of two parts: The first part describes characteristic properties of roads in the real world and in aerial imagery, and represents a road model derived from these properties. The second part defines different local contexts and assigns those to the global contexts.

3.2 Road Model

Besides an (at least partially) homogeneous surface and more or less densely arranged vehicles, one obvious feature of roads in urban areas are road markings separating a road into different lanes. To make use of them, we extended the model of our previous work and now model roads and complex junctions as a combination of several lanes consisting of one or more lane segments. Dashed or solid linear markings define the border of a lane segment. The road model condensed from the findings of the previous section is illustrated in Fig. 2 in form of a hierarchical semantic network.

The model 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 and their relations. On this level the road network consists of junctions and road links that connect junctions. Road links are constructed from road segments. In fine scale, road segments are aggregated by lanes, which consist of pavement and markings. For markings there are two specializations: Symbols and line-shaped markings. The concepts of the real world are connected to the concepts of the *geometry and material* level via *concrete* relations (Tönjes, 1997), which connect concepts representing the same object on different levels. The geometry and material level is an intermediate level which represents the 3D-shape of an object as well as its material (Clément et al., 1993). The idea behind this level is that in contrast to the *image* level it describes objects independently from sensor characteristics and viewpoint. Road segments are linked to the "straight bright lines" of the image level in coarse scale. In contrast to this, the pavement as a part of a road segment in fine scale is linked to the "elongated bright region" of the image level via the "elongated, flat concrete or asphalt region."




Figure 1: Roads in urban areas

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, from fine scale are integrated. In this way the extraction benefits from both scales.

3.3 Context Model

The road model presented above comprises knowledge about radiometric, geometric, and topological characteristics of roads. This model is extended by knowledge about context: So-called *context objects*, i.e. background objects like buildings, trees, or vehicles, can support road extraction, but they can also interfere (see the discussion of Sect. 3.1). Also, external GIS data can be regarded as context object. Experience has shown that modeling this interaction between road objects and context objects on a local level as well as a global level is an aid for guiding the extraction since the interpretation problem is split into smaller sub-problems which can be solved more efficiently by using specific models and extraction strategies.

In order to capture the varying appearance of roads globally, we distinguish between the *context regions urban, forest, and rural* (cf. (Baumgartner et al., 1997)). Furthermore, we model the local context with so-called *context relations*, i.e., certain relations between a small number of road and context objects. For typical context relations in the rural context region, we refer the reader to (Baumgartner et al., 1997). In the following we turn our focus on context relations in the urban context region (see Fig. 3).

Almost every building in the real world is connected with the road network. The denser the settlement is, the closer the buildings move to the road, and the more parallel is their outline with the road sides. Therefore, this context relation gets more useful for the extraction in downtown areas, where, in some extreme cases, roads and junctions are purely defined by the building outlines. Vice-versa, buildings or other high objects standing close to the road potentially occlude larger parts of it or cast shadows on it. Hence, a context relation "occlusion", gives rise to select another image providing a better view on that particular part of the scene, whereas a context relation "shadow" can tell an extraction algorithm to choose modified parameter settings.

Both context relations imply that roads lie beneath the surrounding objects. Consequently, there is no need to search for roads on locally high objects.

For some settlements, road axes might be available digitally. Such kind of information can be integrated in a very consistent way by using a context relation that models parallelism and closeness between an extracted piece of road and the mapped road axis. With such a context relation, cues are provided where a road *might* be present. Nevertheless, the extraction has to prove independently *if* the road truly exists.

Vehicles are related with a road, or more specifically, with a lane segment by means of occluding the lane's pavement. However, since vehicles drive or stand collinear with the a lane (at least in most cases), we can directly use a detected vehicle or vehicle convoy for road extraction – in particular, we treat a detected vehicle as lane segment. By doing so, we need not to take care of moving vehicles if we want to fuse extraction results achieved from images taken at a different time.

In addition to relations between road and context objects, we also consider relations between the object and its sub-structures. This is best exemplified with orthogonal markings at the end of a lane. In most cases, they define the end of a lane and relate it to a junction. Fig. 3 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.

Note, however, that the use of knowledge about local context and the verification of specific relations between local objects

Figure 2: Road model

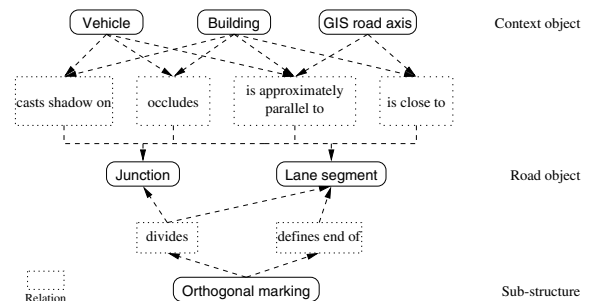
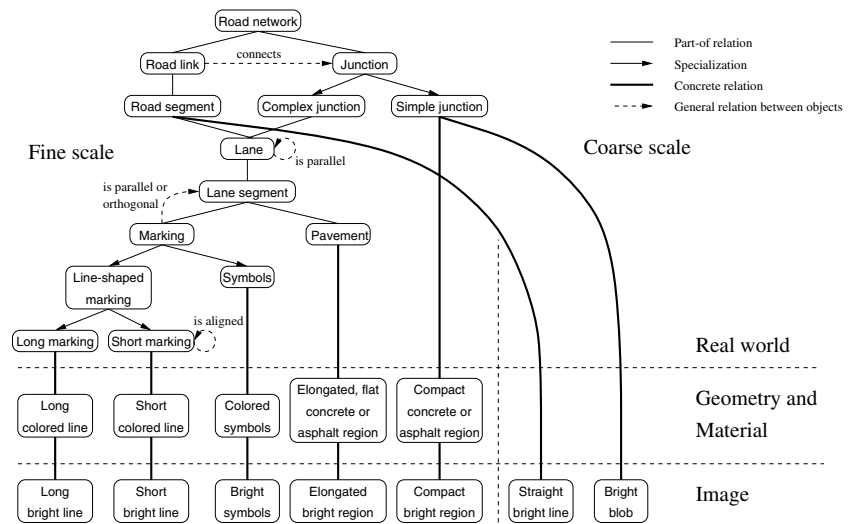


Figure 3: Context relations for urban areas

will in most cases be possible in high resolution imagery only, because the image features which contribute to the local context are usually not very prominent. Therefore, the local context is more tightly connected with the high resolution, whereas information about global context usually can be derived from images with a resolution > 2 m and is useful to guide the road extraction in both scales.

4 EXTRACTION STRATEGY

4.1 Overview

Generally speaking, the extraction strategy inheres knowledge how and when certain parts of the road and context model are optimally exploited. In some reasonable cases, this knowledge is easy to implement as a set of predefined fixed rules, e.g., a road never runs through a house (at least in our model). However, the flexibility of rule-based systems is well-known to be rather limited. Dynamic control systems, e.g., based on Bayesian networks overcome such a drawback. However, they can get computationally expensive rather quickly. Therefore, we plan to implement a control strategy which applies dynamic decision making to smaller sub-problems, e.g., finding reasonable parameter settings for the extraction of markings in shadow regions of a particular scene, whereas the overall flow of the extraction, i.e., (1) Context-based data analysis, (2) Extraction of salient road segments, and (3) Road network completion is defined as shown in Fig. 4.

Since the road model incorporates 3D information as well as small sub-structures to a considerable extent, the extraction relies on one side on aerial imagery consisting of overlapping gray scale images with a fairly high resolution (< 15 cm) and on the other side on an accurate Digital Surface Model (DSM), e.g., as it can be derived from dense airborne laser data. In contrast to other approaches, we neither use orthophotos for extraction nor we extract completely in 3D as, e.g., (Gruen and Li, 1997) do, by using multiple images simultaneously. The latter procedure is conceptually elegant, but it inherently implies matching procedures throughout the extraction process, which become increasingly burdensome with higher scene complexities. Furthermore, we want to avoid feature extraction in orthophotos – despite of using accurate DSM information. Inaccuracies of the DSM due to erroneous height measurements, filtering, resampling, or moving objects still remain and, for instance, they would disturb useful collinear properties of image structures like road markings. Hence, during extraction, we separate image and height information to a certain extent:

At the first stage, down-sampled *image information* is used to extract the global context regions (see (Baumgartner et al., 1997)). Contrarily, we use *DSM information* to analyze height dependent context relations, e.g., for the prediction of shadow regions and occlusions. In dense built-up areas approximate road axes and road sides are derived from the DSM, whereas, e.g. in suburban areas, the image information is used for this. Thereafter, we project the results of context analysis in each image separately. This enables us to treat basic bottom-up procedures purely as 2D-problems, e.g., the extraction and grouping of linear features for hypothesizing vehicle outlines and for detecting groups of markings.

At the next stage, if enough information is collected to construct road objects of a certain spatial extent – in particular, lane segments – we include 3D-information by back-projecting the individual objects on the height model. Here, we fuse the results of the individual images in order to achieve a consistent data set. Optionally, results of road extraction in different context regions can be integrated at this point. Once the lane segments are fused and aggregated to more complex objects (lanes and finally road segments), connection hypotheses between the road segments can be generated. Experience from road extraction in the rural context region has shown that this should be done both locally and globally. Thereafter, verification is carried out in the original images using context relations and standard verification techniques like snakes and homogeneity tracking. The accepted hypotheses are fused with the already

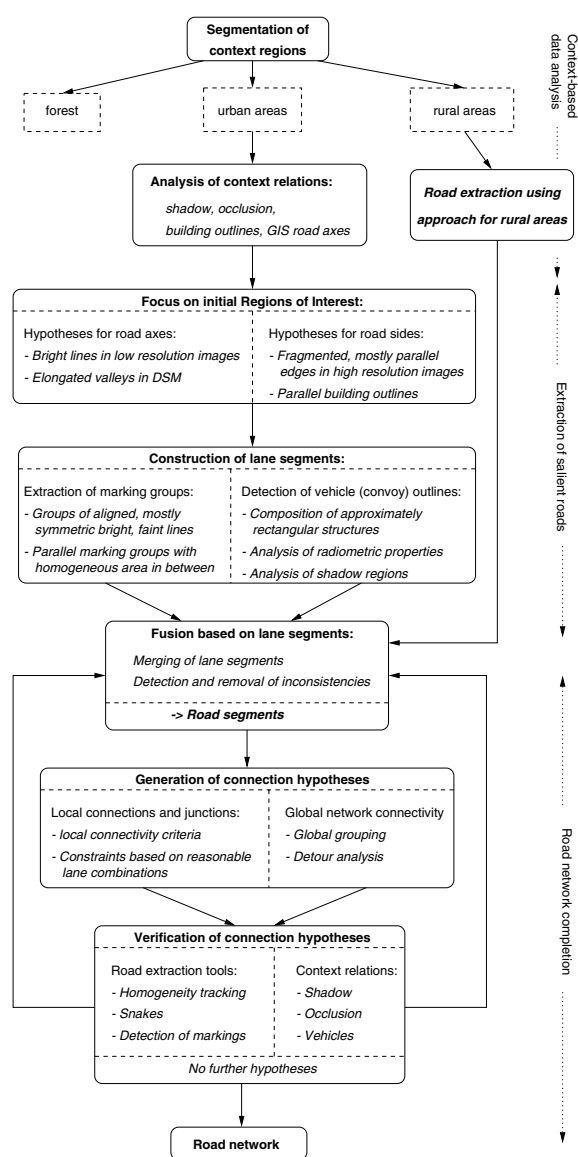


Figure 4: Extraction Strategy

extracted roads, again, and new connections are hypothesized. This process iterates until no more connections are found to be correct.

4.2 Preliminary results

In the first phase of realization of the proposed concept, we focused on one hand on the context-based data analysis, in particular, the segmentation of Regions of Interest (RoI) and the delineation of approximate position and direction of potential roads, since context analysis plays a key role to reduce the scene complexity. On the other hand we implemented modules for the construction of lane segments based on the extraction and grouping of markings. For vehicle detection, some preliminary steps have been undertaken. By this, the contribution of small sub-structures and context objects to road extraction could be investigated. In the following, an example is described which shows results that can be achieved with the modules implemented up to now. For implementation issues, we refer the reader to (Hinz et al., 1999).

The first example shows imagery and DSM of the downtown area of Washington (DC) (~ 20 cm, resp. 3 m resolution). RoIs are segmented using the context relation that most buildings are higher than the road surface. Therefore, the parts that correspond to locally high regions in a DSM are removed from the image. The segmentation procedure compares a smoothed version of the DSM with the original DSM and removes regions where the height difference between both data exceeds a threshold. Furthermore, approximately parallel building outlines are detected by searching for elongated valleys in the DSM. Figure 5 shows the down-sampled image, the DSM image, and the segmented image with the detected valleys. The results are then transformed to the original image resolution followed by the extraction of thin, bright lines. (see Fig. 6).

Thereafter, an iterative graph-based grouping algorithm is applied to group the lines into extended linear objects according to perceptual principles: absolute and relative proximity of lines as well as their continuation (see Fig. 7 a). In regions, where the context analysis was able to find hypothetical road center axes from parallel building outlines, i.e., the valleys in the DSM, only marking groups are kept that show good parallelism with DSM-valleys. The grouping procedure results in a set of unique and topologically consistent marking groups, from which hypotheses for lane segments are generated. We first find parallel marking groups and define their medial axis as the lane axis. However, due to occluding vegetation or parking cars, we can rarely detect road markings at the road sides, even if they are painted there. Hence, we construct additional hypotheses for lane segments on those portions of each side of a particular marking group where no parallel relation to another group could be established.

Finally, the hypotheses are validated: The surface of a lane segment should be homogeneous in the direction of a lane. In road regions where this criterion is not fulfilled, a car must be present. We therefore check the radiometry of lanes in one pixel wide stripes along the respective lane axis. To this end, we shift a one dimensional mask over each stripe and calculate mean brightness and variance. Please note, that the lanes are not constrained to be straight, though some generic knowledge about the geometry of lanes by means of lower bounds for their length and curvature is included (see, e.g., the upper left corner of Fig. 7 b).

Homogeneous parts of a lane segment are labeled as "good" hypotheses, whereas larger portions with high variance indicate abnormal changes in the surface and are labeled as search regions for vehicles. Figure 7 b) visualizes selected bright and homogeneous lane segments, but no hypothesis is ultimately rejected at this point. Some of the

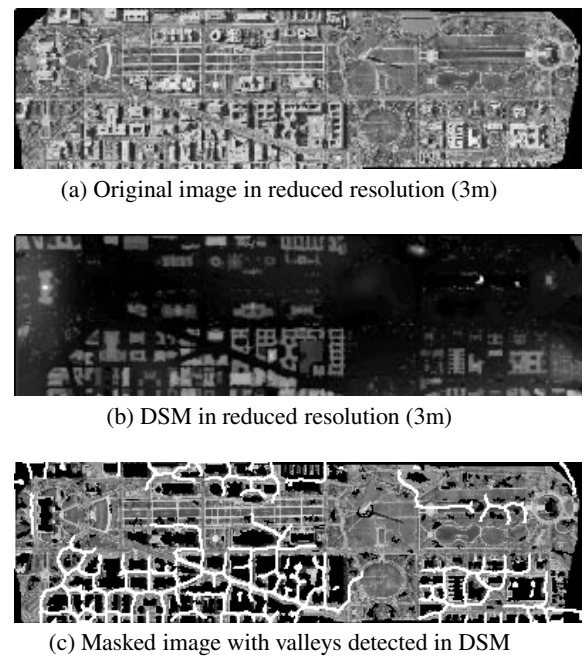


Figure 5: Context analysis with DSM




Figure 6: (a) DSM-valleys projected on image patch, (b) result of line extraction

hypotheses are not selected because of dark shadows and some of the correct hypotheses are interrupted because of cars. Dark regions can be detected by shadow prediction and verification using the respective context relation. Since we already use this kind of information in our road extraction system for rural areas, extracting the exact shadow region in urban areas was not of our primary interest. We rather turned the focus of our work on the detection of vehicles which should provide the system an explanation why such inhomogeneities inside of lane segments occur.

We currently work on the implementation of a model-based scheme for vehicle detection. It starts with the extraction of rectangular edge structures. Then, adjacent rectangles are iteratively concatenated to set up 2D hypotheses for the outlines and center-lines of isolated vehicles or vehicle convoys. Basically, this restricts the hypothesis formation to approximate nadir views. However, since we generally have to select such views in urban areas due to occluding objects, this seems to us being no true limitation. Figure 8 a) shows a particular part of a lane axis (in black) which was not accepted by the homogeneity check. (In fact, it is the horizontal road, resp. lane, in the center of the image shown in 7 b)). The result of hypothesis formation is visualized in Fig 8 b), extracted edge segments are plotted in black. The generated car hypothesis (shown in white) consists of individual rectangles which complete the fragmented edges resulting in a chain of rectangles. The medial axis of the rectangle chain corresponds to the center-line of the hypothesized vehicle.

Once, we have generated a hypothesis for a vehicle's center-line, we are able to check for significant geometric and radiometric symmetries along and across the vehicle. In aerial imagery, most vehicles are characterized by dark windows. Furthermore, their front, top, and rear, show similar reflectance. The use of color images could be a potential aid for the decision which rectangles belong to one single car and how to discriminate different cars within a convoy. With this information we should be able to select a few, rather generic 3D models from a database containing a limited set of vehicle models. Verification is carried out by projecting the selected models back to the image and matching, e.g., their wire-frame representation to the grayvalue gradients. A good match between the model and the image indicates that the system has selected a reasonable model. An additional and, in fact, strong verification for a vehicle is its shadow region. By illuminating the 3D model we can predict the corresponding shadow region on the road surface and search in the image for the respective features (i.e., dark region, corresponding edges, dark side of the edges against the sun)

However, it should be noted that our primary goal is the extraction of lanes and not the detection of individual vehicles. It is sufficient for this purpose to know *that* a vehicle occludes the road, but not *what kind* of vehicle it is.

5 SUMMARY AND OUTLOOK

We presented our concept for automatic road extraction in urban areas. It is based on a detailed model for roads, including lanes and road markings, and their context. By using image and DSM information, the extraction strategy first exploits global context knowledge, in order to reduce the inherently high complexity of urban scenes. Then, it focuses on image parts where initial lane and road hypotheses are rather easy to detect. Iteratively hypothesizing and verifying connections between already extracted road segments completes the network. So-called context relations help to analyze and explain abnormal changes in the expected appearance of the road. Especially, the context relation that models vehicles on roads is of great importance. The results of the currently implemented modules encourage us to further realize this concept. However, there are still many steps to go and there are still many questions to be answered.

REFERENCES

Airault, S., Ruskoné, R. and Jamet, O., 1994. Road detection from aerial images: a cooperation between local and global methods. In: J. Desachy (ed.), *Image and Signal Processing for Remote Sensing*, Proc. SPIE 2315, pp. 508–518.




Figure 7: (a) Detected groups of markings, (b) lane segments with bright homogeneous interior




Figure 8: Rectangular car hypothesis

- Baumgartner, A., Eckstein, W., Mayer, H., Heipke, C. and Ebner, H., 1997. Context-Supported Road Extraction. In: (Gruen et al., 1997), pp. 299–308.
- 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.
- Bogenberger, K., Ernhofer, O. and Schütte, C., 1999. Effects of Telematic Applications for a High Capacity Ring Road in Munich. In: *Proceedings of the 6th World Congress on Intelligent Transportation Systems*, Toronto.
- 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.
- de Gunst, M. E., 1996. Knowledge-based Interpretation of Aerial Images for Updating of Road Maps. PhD thesis, Delft University of Technology, Delft.
- Englisch, A. and Heipke, C., 1998. Erfassung und Aktualisierung topographischer Geo-daten mit Hilfe analoger und digitaler Luftbilder. *Photogrammetrie - Fernerkundung - Geoinformation (PFG)*.
- Förstner, W., 1999. 3D-City Models: Automatic and Semiautomatic Acquisition Techniques. In: D. Fritsch and R. Spiller (eds), *Photogrammetric week'99*, Wichmann.
- Fuchs, C., Gülch, E. and Förstner, W., 1998. OEEPE Survey on 3D-City Models. OEEPE Publication, 35.
- Geman, D. and Jedynek, B., 1996. An active testing model for tracking roads in satellite images. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 18(1), pp. 1–14.
- Gruen, A. and Li, H., 1997. Linear feature extraction with LSB-snakes. In: (Gruen et al., 1997), pp. 287–298.
- Gruen, A., Baltsavias, E. and Henricsson, O. (eds), 1997. *Automatic Extraction of Man-Made Objects from Aerial and Space Images (II)*. Birkhäuser Verlag, Basel.
- Harvey, W. A., 1999. Performance Evaluation for Road Extraction. In: *International Workshop on 3D Geospatial Data Production: Meeting Application Requirements*, Paris.
- Heller, A., Fischler, M., Bolles, R. and Connolly, C., 1998. An Integrated Feasibility Demonstration for Automatic Population of Spatial Databases. In: *Image Understanding Workshop*.
- Hinz, S., Baumgartner, A., Steger, C., Mayer, H., Eckstein, W., Ebner, H. and Radig, B., 1999. Road extraction in rural and urban areas. In: W. Förstner, C.-E. Liedke and J. Bückner (eds), *SMATI'99: Semantic Modeling for the Acquisition of Topographic Information from Images and Maps*, pp. 133–153.
- Mayer, H., 1999. Automatic Object Extraction from Aerial Imagery – A Survey Focussing on Buildings. *Computer Vision and Image Understanding* 74(2), pp. 138–149.
- 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.
- Mayer, H., Laptev, I. and Baumgartner, A., 1998. Multi-Scale and Snakes for Automatic Road Extraction. In: *Fifth European Conference on Computer Vision*, pp. 720–733.
- Price, K., 1999. Road Grid Extraction and Verification. In: *International Archives of Photogrammetry and Remote Sensing*, Vol. XXXII, part 3-2W5, International Society for Photogrammetry and Remote Sensing.
- Ruskoné, R., 1996. Road Network Automatic Extraction by Local Context Interpretation: application to the production of cartographic data. PhD thesis, Université Marne-La-Vallee.
- Ruskoné, R., Airault, S. and Jamet, O., 1994. Road network interpretation: A topological hypothesis driven system. In: *International Archives of Photogrammetry and Remote Sensing*, Vol. XXX, part 3/2, pp. 711–717.
- Ruskoné, R., Guiges, L., Airault, S. and Jamet, O., 1996. Vehicle detection on aerial images: A structural approach. In: *13th International Conference on Pattern Recognition*, Vol. III, pp. 900–903.
- Steger, C., 1998. An Unbiased Detector of Curvilinear Structures. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 20(2), pp. 113–125.
- Tönjes, R., 1997. 3D Reconstruction of Objects from Aerial Images Using a GIS. In: *International Archives of Photogrammetry and Remote Sensing*, Vol. 32, Part 3-2W3, pp. 140–147.
- Trinder, J. and Wang, Y., 1998. Knowledge-Based Road Interpretation in Aerial Images. In: *International Archives of Photogrammetry and Remote Sensing*, Vol. XXXII, part 4, pp. 635–640.
- Vosselman, G. and de Gunst, M., 1997. Updating Road Maps by Contextual Reasoning. In: (Gruen et al., 1997), pp. 267–276.
- Wiedemann, C., 1999. Automatic completion of road networks. In: *Proceedings of the ISPRS Joint Workshop 'Sensors and Mapping from Space 1999'*, Hanover.
- Wiedemann, C. and Hinz, S., 1999. Automatic Extraction and Evaluation of Road Networks from Satellite Imagery. In: *International Archives of Photogrammetry and Remote Sensing*, Vol. XXXII, part 3-2W5, International Society for Photogrammetry and Remote Sensing, pp. 95–100.
- Zhang, C. and Baltsavias, E., 1999. Road Network Detection by Mathematical Morphology. In: *International Workshop on 3D Geospatial Data Production: Meeting Application Requirements*, Paris.