

Semantic Objects and Context for Finding Roads

A. Baumgartner^a, C. Steger^b, H. Mayer^a, W. Eckstein^b

^aLehrstuhl für Photogrammetrie und Fernerkundung
Technische Universität München
Arcisstr. 21 D-80333 München Germany

^bForschungsgruppe Bildverstehen (FG BV), Informatik IX
Technische Universität München
Orleansstr. 34 D-81667 München Germany

ABSTRACT

This paper presents a multi-resolution approach for automatic extraction of roads from digital aerial imagery. Roads are modeled as a network of intersections and links between the intersections. For different context regions, i.e., rural, forest, and urban areas, the model describes different relations between background objects, e.g., buildings or trees, and semantic road objects, e.g., road-parts, road-segments, road-links, and intersections. The classification of the image into context regions is done by texture analysis. The approach to detect roads is based on the extraction of edges in a high resolution image and the extraction of lines in an image of reduced resolution. Using both resolution levels and explicit knowledge about roads, hypotheses for roadsides are generated. The roadsides are used to construct quadrilaterals representing road-parts and polygons representing intersections. Neighboring road-parts are chained to road-segments. Road-links, i.e., the roads between two intersections, are built by grouping of road-segments and closing of gaps between road-segments. Road-links are constructed using knowledge about context.

Keywords: Road Extraction, Multi-Resolution, Context

1. INTRODUCTION

In digital photogrammetry semantic tasks, like data capture and update of geographic information systems (GIS), still have to be performed manually by a human operator. This is very time consuming and expensive. Therefore, automatic solutions are highly desirable. Even though for GIS update, i.e., verification and acquisition, old GIS data can make object extraction easier, the proposed approach is restricted to object acquisition (here: roads) without using prior GIS information (for work in this area see, e.g., Refs. 1,2). This is done because automatic acquisition shows the potential and limits of an extraction scheme and can deepen the understanding of image interpretation. What is more, the extraction of new roads is needed for the update in any case.

A lot of work exists on the extraction of roads. The most common methods are extraction or tracking of lines in small scale imagery and profile matching or detection of roadsides in large scale images. The approaches vary in the way how different methods are combined as well as how additional knowledge, e.g., geometrical constraints, are incorporated. A main criterion to distinguish the works is the interaction of a human operator. In semi-automatic schemes an operator selects an initial point and a direction for a road tracking algorithm (cf. Refs. 3–6). In Refs. 7–9 the operator marks a few points of a road segment and an algorithm, e.g., based on dynamic programming, finds the road which connects these points (also in 3D for more than one image). This is advantageous because the path of the road is more constrained, and a more reliable handling of obstacles is possible. A similar approach based on so-called “ziplock” snakes is given in Ref. 10. These semi-automatic approaches can be extended to fully automatic operation by means of automatic seed point detection (cf. Ref. 11). In Ref. 12 seed points for the extraction of the road network are centers of elongated regions found by a watershed-based segmentation on a gradient image. Based on the elongated regions and their directions, road segments are extracted using the homogeneity of the road

Correspondences:

^aTel: +49-(89)-289 22688; Fax: +49-(89)-2809573; E-mail: {albert|helmut}@photo.verm.tu-muenchen.de

^bTel: +49-(89)-48095 211; Fax: +49-(89)-48095 112; E-mail: {steger|eckstein}@informatik.tu-muenchen.de

surface. In order to extract the road network geometrical constraints are taken into account, and hypotheses about connections between road segments are checked.

If relations between roads and other objects, like cars, buildings, or trees, are neglected a reliable extraction is often difficult. Background objects can have a strong influence on the characteristics of roads, or at least on the appearance of roads in aerial imagery, e.g., high objects causing occlusions or casting shadows. Therefore, Ref. 13 strongly demands the use of contextual information to guide extraction algorithms for complex scenes. In Ref. 14 the detection and grouping of cars is used for road extraction in urban areas where profile matching or detection of roadsides would probably fail. Ref. 15 analyzes the influence of context on the accuracy of roads in a cartographic database.

As distinct characteristics of roads can be detected best at different resolution levels and in different contexts this paper proposes a multi-resolution extraction scheme integrating context information. On a global level three kind of so-called *context regions* are introduced: urban, rural, and forest. For each context region special relations between roads and background objects are considered. The model in Section 2 describes the road network as a hierarchy of road objects. The relations between road objects and background objects are described by so-called *context sketches*. In contrast to the context regions context sketches are used to model contextual knowledge on a local level. Section 3 shows the fusion of different resolution levels for the extraction of roadsides and the construction of road objects, along with their attributes and methods. The segmentation of the image into different context regions based on texture analysis, and how the context information is used, is explained in Section 4. Finally, in Section 5 conclusions are given.

2. MODEL

In order to extract roads it is necessary to have a clear idea about the concept *road*. Hence, a model should comprise explicit knowledge about road width, parallelism of roadsides (geometry), reflectance characteristics (radiometry), network structure (topology), and relations to other objects like buildings or trees (context). The model described below is split into two parts. One part describes characteristics of roads (Sect. 2.1). The other part defines relations to background objects, called context sketches here, and assigns these relations to different context regions (Sect. 2.2). In this way the complex model of the object *road* is divided into more specific sub-models which are adapted to contextual environment (urban, suburb, rural, forest), sensor, and scale (resolution). The sub-models emphasize certain characteristics of the objects and therefore can be regarded as specialized models.

2.1. Model of Road Objects

Roads in Real World: In the real world, the characteristics of roads can be described as consequence of their function for human beings.¹⁶ Roads make space accessible and therefore exhibit the form of a network connecting all areas inhabited and used by human beings. The denser an area is inhabited and the more intensive it is used, the denser is the road network of the area. Corresponding to their importance, network components are classified into a hierarchy of different categories with different attributes. Roads of national importance, i.e., highways, are much wider than roads accessing agricultural or forest areas. Their relation with the terrain surface differs according to their function and category. Agricultural paths and less important feeder roads are more closely following the natural terrain surface than highways do, the latter ones serving as fast connections between densely settled areas. Roads of different categories therefore differ with respect to minimum curvature radius and maximally allowed slope. As a consequence of greater curvature radius and more gentle slope it happens more often that roads of a greater importance run along embankments, over bridges or in tunnels. On closer inspection, distinct parts of roads can be attributed by the type of surface material, the existence and type of road markings, or the existence of sidewalks and cycle-tracks.

Roads in Aerial Images: The appearance of roads in digital imagery strongly depends on the sensor's spectral sensitivity and its resolution in object space. Usually roads show up with specific grey-values, especially in color and infra-red images. The remainder of this paper is restricted to gray-scale images and only scale dependencies are considered. In images with varying resolution different characteristics of roads are appearing. In images with low resolution, i.e., more than 2 m per pixel, roads mainly appear as lines establishing a more or less dense network. In images with a higher resolution, i.e., less than 0.50 m, roads are depicted as elongated homogeneous areas with more or less parallel borders and almost constant width. Usually the brightness of roads is different from the surroundings

and therefore the roadsides can be extracted using edge detectors. In the higher resolution the attainable geometric accuracy is better, but background objects like cars, trees, and buildings disturb the road extraction to a higher degree.

Road Model: From the last paragraph follows that fusion of low and high resolution results can contribute to improve the reliability of the road hypotheses. Additionally, details like road markings which can be recognized at a resolution of less than 0.20 m, can be used as evidence to corroborate the detected road hypotheses. On the one hand, using multiple resolution levels improves robustness of road extraction. On the other hand, it results in the necessity to use different features at each resolution level, and to simultaneously combine all features of all resolution levels into one road model. The semantic network in Figure 1 illustrates the road model condensed from this.

The model is split into three levels, defining different points of view. The *real world* level consists of the objects and their relations on a natural language level. The road of this level consists of intersections and road-links, which connect intersections. A road-link is assembled from short parts, the road-segments. In large scale a road-segment is constructed of road-parts (for the motivation of their use see Sect. 3.2) which in turn comprise the pavement and the markings (solid or dashed). The objects in the *real world* level are connected to the objects in the *geometry and material* level by means of the *concrete* relation which connects *concepts* describing the same object on different levels, i.e., from different points of view.¹⁷ The *geometry and material* level is an intermediate level which represents the three-dimensional shapes of objects as well as their material.¹⁸ This level has the advantage that it represents objects independently from sensor characteristics and viewpoint, which is in contrast to the *image* level.




Figure 1. Road model: different resolutions and points of view

Road-segments in small scale are linked to a mostly straight bright line in the *image* level. In contrast to this, the pavement of the large scale is linked to the elongated bright area in the *image* level via the elongated flat concrete or asphalt area in the *geometry and material* level. The markings are related to bright lines via colored lines.

Whereas the large scale gives detailed information, the small scale adds global information. If the information of both levels is fused, false hypotheses for roads are eliminated by using the abstract small scale information, while integrating details from the large scale (e.g., the correct width of the roads). With this, the advantages of both scales are merged.

2.2. Context

The road model of the last section comprises knowledge about radiometric, geometric, and topological characteristics. But besides features which refer to the road itself, it is also necessary to consider relations between roads and other

objects like buildings, trees, and cars. This is, because these background objects on the one hand support the evidence of roads and on the other hand hinder road extraction.

Context Sketches: The notion “context sketch” is introduced to describe typical relations between road objects and background objects. The context sketch *building_driveway_road-segment* (cf. Fig. 2) models the relations between a road-segment, a driveway (i.e., a small and short road-part), and a building at the end of the driveway. The driveway is often almost perpendicular to the road-segment and the building is in many cases parallel to the road-segment. The context sketch *occlusion_shadow* consists of a hypothetical road-part which connects two road-segments and a high object next to the hypothetical road-part. Figure 3a) illustrates an occlusion and Figure 3b) a shadow cast on the road by high object. Another context sketch describes the relation between road-segments and driveways to agricultural fields (*rural_driveway*). Figure 4a) explains the situation where a driveway causes a gap between two road-segments and only one roadside is found. The context sketch *vehicle_road* describes a dark or bright blob on a road-part or on a road-segment. To be a vehicle the blob must have a certain size and it should not be centered on the road (cf. Fig. 4b). The relation between roads and parallel objects, like sidewalks and cycle-tracks, is defined by the context sketch *parallel_object*.




Figure 2. Context sketch *building_driveway_road-segment*




Figure 3. Context sketch *occlusion_shadow*




Figure 4. Context sketches: a) *rural_driveway* b) *vehicle_on_road*

These basic context sketches are modules that can be aggregated into more complex context sketches depicting, e.g., the interaction between *building_driveway_road-segment* and *occlusion_shadow* or the interaction between

building_driveway_road-segment and *vehicle_on_road*.

Context Regions: Not every context sketch always has to be taken into account. This is in contrast to approaches representing the whole scene with one model.¹⁹ The relevance of features and relations depends also on the global context. Roads in urban or suburban areas look quite different and have other relations than roads in rural or forest areas. Therefore, this paper proposes to use different features and relations, i.e., context sketches, not only at multiple resolution levels but also within so-called context regions, for which *suburb_urban*, *forest*, and *open_rural* areas are distinguished here. Information from a GIS or a (pre-)segmentation of the image into these regions provides global à priori information about the characteristic features and their relations. For example, buildings in downtown areas are — in contrast to buildings in rural areas — very close and highly parallel to roads; sidewalks and cycle-tracks are more likely to appear in urban areas; in rural areas single trees and single buildings might hinder extraction, whereas in forest regions mainly shadows and occlusions pose problems. This information about existence or proximity of background objects makes it possible to choose the most appropriate extraction algorithm and to determine the meaning of distinct road-parts and road-segments. The assignment of the context sketches to the context regions is shown in Figure 5.

Figure 5. Context regions and context sketches for roads

3. STRATEGY

The main principle in the proposed extraction scheme is to guide road extraction by already available information. Therefore, the most salient objects are extracted first. Depending on the area where roads are to be extracted, different features are easiest to detect. In urban and forest areas radiometric and geometric knowledge alone is not sufficient due to shadows and occlusions. On the other hand, with a simple model, relying only on attributes of the road itself, good results are expected for rural areas. According to these considerations the road extraction starts in rural areas mainly based on low level algorithms. The intermediate results after this step can then serve to guide road extraction in urban and forest areas. To enable the reasoning in object space (i.e., to use information about the scene in meters and not in pixels), the road extraction is done in an orthophoto. This gives enough information in most cases, since roads are planar objects. According to the considerations explained above, not only images at original resolution but also images at reduced resolution are used. A segmentation of the image into different context regions is achieved by texture analysis in the reduced image.

In a first step, roadsides are extracted using a multi-resolution approach, which integrates edge extraction in a high resolution image and line extraction in a low resolution image (Sect. 3.1.). In a second step, road objects, like road-parts, road-segments, and intersections, are constructed from the roadsides (Sect. 3.2). The road objects have attributes and methods to enable a more sophisticated examination of reasons for gaps between road-parts and road-segments. The role of context information for road extraction is shown in Section 4.

3.1. Extraction of Roadsides

According to the radiometric part of the road model roadsides correspond to edges in high resolution images (here ca. 0.25 m). As candidates for roadsides those edges are selected which are parallel and border homogeneous regions. However, extraction at high resolution is error-prone due to disturbing objects and noise. Therefore line extraction in an image of reduced resolution (ca. 2 m) is integrated to get less erroneous hypotheses. At the low resolution level the global network structure of the roads can be seen, and small disturbances, like cars on the road or shadows cast by trees, do not pose as many problems as at high resolution level due to the smoothing of the image (cf. Fig. 6). A drawback of the low resolution is that the geometrical accuracy is poor compared to high resolution.




Figure 6. a) Image at low resolution. b) Hypotheses for road axes at low resolution.



Figure 7. a) Input to the fusion process: hypotheses for road axes (dashed, black), hypotheses for roadsides (solid, white) and other edges (dotted, white). b) Final roadside hypotheses.

As motivated in Section 2.1, the results should be fused, to make optimal use of the advantages of both resolution levels. Input for the fusion are hypotheses for road axes from low resolution, hypotheses for roadsides from high resolution, and the original edge polygons. The basic strategy is to select the results of both levels supporting each other. Hypotheses for roadsides which enclose a homogeneous area and have a hypothesis for a road axis in-between serve as initial hypotheses. Details about this extraction of roadsides using different resolution levels are given in Ref. 20. Figure 7a) shows the input to the fusion process. The final roadside hypotheses are presented in Figure 7b).

3.2. Extraction of Semantic Objects

The results obtained in the last section can be regarded as mid-level segmentation results. This means, that even if they look like roads in Figure 7b), they have not been extracted so far; candidates for roadsides have been found. In order to facilitate further processing, semantically meaningful objects with rich attributes have to be generated.

As shown in Figure 1, the road network consists of *intersections* and *road-links*, which connect two intersections, and can therefore be regarded as a graph. However, it is rarely possible to extract a single connected road-link between two intersections because of occlusions and other problems. Therefore, the concept *road-segment* is used to describe a single, connected piece of road detectable in the image. At low resolution this concept corresponds directly to the bright lines extracted at low resolution. At high resolution a road-segment can be described by its enclosing polygon. However, it is quite difficult to generate this polygon directly in a topologically correct way. Therefore, *road-parts* are introduced as intermediate concepts. They have a simple geometric description, namely the minimal quadrilateral built by two parallel and overlapping roadsides. Road-segments can then simply be created by aggregating adjacent road-parts. Similarly, at low resolutions intersections are junctions of more than two lines. Conversely, at high resolution they are regions described by their enclosing polygon.




Figure 8. Construction of road-parts and road-segment




Figure 9. a) Road-parts b) Road-segments

Extraction of road-parts: The first step to generate the different types of road objects is to generate road-parts

from the results obtained in Section 3.1. This is done by splitting parallel roadsides into non-overlapping quadrilaterals by projecting each point of the left and right parallel polygons perpendicularly onto the other polygon, as shown in Figure 8. Furthermore, additional roadsides found by the fusion step are checked whether they are parallel, overlapping, and close to each other. The same procedure is applied to these roadsides to generate additional road-parts. Each road-part is described by the four points of the quadrilateral, which are oriented clockwise. Furthermore, there are functions in the corresponding class to calculate the center axis of each road-part, and to convert the road-part into a region in the image. The latter point facilitates the selective extraction of features on the road-part, e.g., road markings. Figure 9a) shows the extracted road-parts.

Aggregation of road-parts into road-segments: When the road-parts have been extracted it is straightforward to group them into road-segments by collecting all adjacent road-parts (cf. Fig. 8). Road-segments inherit much of their attributes and functionality from their constituting road-parts, e.g., the abilities to generate their center axis, and to be converted into regions. Figure 9b) shows the road-segments generated from the road-parts in Figure 9a). It is apparent that most of the roads are correctly extracted. It can also be seen that sometimes the generation of road-parts fails to cover the entire area between parallels because at some areas no quadrilaterals can be generated, especially in highly curved areas. There are also some small gaps at places where only one roadside has been found. However, this can easily be corrected. The road-segments thus created can be aggregated into road-links by grouping processes. The bridging of gaps will be described in Section 4.




Figure 10. Construction of intersections

Extraction of intersections. In Figure 1 intersections are areas in the image where more than two road-links meet. At low resolution they are points where three or more lines meet (junctions). At high resolution they are complicated objects defined by their enclosing polygon. Candidates for this polygon are the roadsides remaining after the steps above, under the assumption that they border a road-like region, but are not parallel. However, to sort out which roadsides belong to a particular intersection is not trivial. Therefore, in this step information from the low resolution level is used. Hypotheses for intersections are detected at low resolution by constructing a road network graph and selecting junctions. Since the line detector* fails to extract junctions under predictable circumstances, the algorithm completes the graph by adding junctions at places where a line end is very close and almost perpendicular to another line. The adjacent lines are then examined whether they cross road-segments at high resolution. This procedure yields road-segment candidates for a particular intersection. Since the lines adjacent to a junction can be quite long, and cross several road-segments at high resolution, the closest road-segments on each line are selected

*Details about the line extraction algorithm are given in Ref. 21.



Figure 11. a) Road-segments (white), lines (black) and junctions (white points). b) Polygons of intersection areas (black) and road-segments (white), after closing of small gaps.

as belonging to the intersection, and are used to define a search area for the enclosing polygon of the intersection (cf. Fig. 10). Polygons in this area are then aggregated to the enclosing polygon of the intersection by tracing them from the end points of the adjacent road-segments, and, if necessary, by closing gaps where no roadsides or edges were found. Figure 11a) shows the road-segments and the network graph derived from low resolution, while Figure 11b) displays the extracted intersection polygons. When four or more road-segments form the intersection the procedure has to be slightly modified, since the line extractor will usually extract two closely adjacent junctions of three lines in this case. Therefore, usually only two road-segments will be found for each of the two junctions. This case can be treated by checking if the line for which no road-segment has been found, connects to another junction where also only two segments have been found, and whether this connecting line is not too long to be on an intersection. Like road-segments, intersections have a rich set of attributes (e.g., the adjacent road-segments) and functions (e.g., they can be converted into discrete regions for low level processing).

4. ROLE OF CONTEXT

4.1. Context Regions

As described in Section 2.2, context regions provide additional global information enabling a more efficient use of knowledge about roads. The internal control of the road extraction scheme depends strongly on the global information provided by context regions. A segmentation of the aerial image into the context regions can be obtained by given GIS data or by texture analysis. Figure 12 shows an example of a texture based segmentation of an aerial image into *open_rural*, *forest*, and *suburb_urban* areas. Texture energy computed with different filter-masks proposed by Ref. 22 at a resolution of 4 m is used as criterion. Ideally results of the texture analysis are combined with available GIS data. The context regions derived from the GIS can be compared with the segmentation results. If the borders of GIS regions and segmented context regions do not differ too much the GIS borders are supposed to be correct otherwise the borders of context regions are taken from the segmentation. However, it is important to note that the GIS borders are not enough, as the goal is update and therefore the GIS can be outdated.

4.2. Road Extraction in *open_rural* Area

This section describes the strategy for the extraction of those parts of the road network which are in the *open_rural* context region. The first step is to construct all road-parts, road-segments and intersections in the *open_rural* area.




Figure 12. Classification at low resolution: a) *open_rural* b) *forest* c) *suburb_urban*

The goal is to close the gaps between the road-segments and to generate road-links, i.e., long road-segments connecting two intersections. Figure 13 shows an example for the input data to this step.




Figure 13. a) Roadside b) Road-segments

In *open_rural* area there are two main reasons for gaps: First, small gaps caused by the segmentation (extraction of roadsides, construction of road-parts, etc.) and second, gaps caused by local disturbances, like single trees next to the road or branching driveways to agricultural land. Gaps of the first kind usually do not exceed more than a few meters, and can be closed if the road-segments are connected by an extracted roadside at one side (cf. Fig. 14a). Gaps of the second kind can be handled using the context sketches *occlusion_shadow* and *rural_driveway*. If a gap can be explained by such a context sketch, a road-part is constructed to connect the two road-segments (cf. Fig. 14b).

After this step, all road-segments ideally should connect two intersections, i.e., they are road-links, or they should end at the border of the *open_rural* area or the image. Road-segments with dead ends inside the *open_rural* area have to be checked whether they are really roads. To accept such road-segments as road-links, an important object like a




Figure 14. a) Road-segments after closing gaps of first-kind b) Road-segments after closing gaps of second-kind (white), Intersection polygons (black)

building has to be close to the dead end.

Additional and very strong evidence for a road is gained from the detection of road markings which appear as weak bright lines in the high resolution image. Of course, reasonable results are possible only because the line extraction is restricted to the hypothetic road area: The markings are very faint, and can only be found using very low thresholds, which would result in a plethora of lines for other parts of the image.

4.3. Outlook: Road Extraction in *suburb_urban* and *forest* Areas

For the extraction of roads in *suburb_urban* and *forest* areas road-segments of the *open_rural* area ending at the border to the other context regions can be used as reliable starting points for a road tracking into the *suburb_urban* and *forest* areas. Hypotheses for paths which connect the starting points are derived from the results of the line extraction at low resolution. The road network is propagated from the border to the center of the *suburb_urban* and *forest* areas. In order to achieve good results with the road extraction at high resolution, it will be necessary to integrate context information, especially about high objects and shadows. The detection of buildings, trees, and shadows is possible using an automatically generated digital surface model (cf. Ref. 23). Also results of approaches for building detection can be used. This additional information will be employed by means of context sketches.

5. SUMMARY

In this paper a multi-resolution approach for road extraction from aerial imagery is presented which makes use of local (sketch) and global (region) context information. The road model describes objects of a road network and relations between these parts and other objects. Specifically roadsides are extracted by fusing of the results of line extraction at low resolution and edge extraction at high resolution. From the result road-parts, road-segments, and intersections are constructed. These road objects have a clear semantics as well as attributes and methods for further reasoning. Integration of context regions restricts possible context sketches (i.e., relations to background objects) and makes use of knowledge more effective.

ACKNOWLEDGMENTS

This work is embedded in the joint research project *Semantic Modeling and Extraction of Spatial Objects from Images and Maps* and funded by the Deutsche Forschungsgemeinschaft (DFG) under grant no. Eb 74-8/2.

REFERENCES

1. M. de Gunst, *Knowledge-Based Interpretation of Aerial Images for Updating of Road Maps*, Netherlands Geodetic Commission, Publications on Geodesy, New Series 44, Delft, The Netherlands, 1996.
2. C. Wiedemann and H. Mayer, "Automatic Verification of Roads in Digital Images Using Profiles," in *Mustererkennung 1996*, pp. 609–618, Springer-Verlag, (Berlin), 1996.
3. D. M. McKeown Jr. and J. L. Denlinger, "Cooperative methods for road tracking in aerial imagery," in *Computer Vision and Pattern Recognition*, pp. 662–672, 1988.
4. S. Airault, R. Ruskoné, and O. Jamet, "Road detection from aerial images: a cooperation between local and global methods," in *Image and Signal Processing for Remote Sensing*, Proc. SPIE 2315, pp. 508–518, 1994.
5. G. Vosselman and J. de Knecht, "Road tracing by profile matching and Kalman filtering," in *Automatic Extraction of Man-Made Objects from Aerial and Space Images*, pp. 265–274, Birkhäuser Verlag Basel, 1995.
6. D. Geman and B. Jedynek, "An Active Testing Model for Tracking Roads in Satellite Images," *IEEE Transactions on Pattern Analysis and Machine Intelligence* **18**(1), pp. 1–14, 1996.
7. A. Gruen and H. Li, "Linear feature extraction with LSB-snakes from multiple images," in *International Archives of Photogrammetry and Remote Sensing*, vol. 31, Part B3, pp. 266–272, 1996.
8. J. C. Trinder, "Extraction of Man-Made Features by 3-D Active Contour Models," in *International Archives of Photogrammetry and Remote Sensing*, vol. 31, Part B3, pp. 874–879, 1996.
9. N. Merlet and J. Zerubia, "New Prospects in Line Detection by Dynamic Programming," *IEEE Transactions on Pattern Analysis and Machine Intelligence* **18**(4), pp. 426–431, 1996.
10. W. Neuenschwander, P. Fua, G. Székely, and O. Kübler, "From ziplock snakes to velcrotm surfaces," in *Automatic Extraction of Man-Made Objects from Aerial and Space Images*, pp. 105–114, Birkhäuser Verlag Basel, 1995.
11. A. Zlotnick and P. Carnine, "Finding Road Seeds in Aerial Images," *Computer Vision, Graphics, and Image Processing: Image Understanding* **57**, pp. 243–260, 1993.
12. R. Ruskoné, S. Airault, and O. Jamet, "Road network interpretation: A topological hypothesis driven system," in *International Archives of Photogrammetry and Remote Sensing*, vol. 30, Part 3/2, pp. 711–717, 1994.
13. T. M. Strat and M. A. Fischler, "The Role of Context in Computer Vision," in *Workshop on Context-Based Vision*, pp. 2–12, IEEE, 1995.
14. R. Ruskoné, L. Guigues, S. Airault, and O. Jamet, "Vehicle detection on aerial images: A structural approach," in *13th International Conference on Pattern Recognition*, vol. III, pp. 900–904, 1996.
15. G. Bordes, G. Giraudon, and O. Jamet, "Road Extraction Guided by a Cartographic Database: Creating a Strategy," in *Workshop on Context-Based Vision*, pp. 30–37, IEEE, 1995.
16. L. Stark and K. Bowyer, "Achieving Generalized Object Recognition through Reasoning about Association of Function to Structure," *IEEE Transactions on Pattern Analysis and Machine Intelligence* **13**(10), pp. 1097–1104, 1991.
17. H. Niemann, G. Sagerer, S. Schröder, and F. Kummert, "ERNEST: A Semantic Network System for Pattern Understanding," *IEEE Transactions on Pattern Analysis and Machine Intelligence* **12**(9), pp. 883–905, 1990.
18. R. Tönjes, "Knowledge Based Modelling of Landscapes," in *International Archives of Photogrammetry and Remote Sensing*, vol. 31, Part B3, pp. 868–873, 1996.
19. T. Matsuyama and V.-S. Hwang, *SIGMA - A Knowledge-Based Aerial Image Understanding System*, Plenum Press, New York, USA, 1990.
20. C. Heipke, C. Steger, and R. Multhammer, "A hierarchical approach to automatic road extraction from aerial imagery," in *Integrating Photogrammetric Techniques with Scene Analysis and Machine Vision II*, D. M. McKeown Jr. and I. J. Dowman, eds., Proc. SPIE 2486, pp. 222–231, 1995.
21. C. Steger, "An Unbiased Detector of Curvilinear Structures," Technical Report FGBV-96-03, Forschungsgruppe Bildverstehen (FG BV), Informatik IX, Technische Universität München, 1996.
22. K. Laws, *Texture Image Segmentation*, Ph.D. thesis, Dept. of Engineering, University of Southern California, 1980.
23. W. Eckstein, "Segmentation and Texture Analysis," in *International Archives of Photogrammetry and Remote Sensing*, vol. 31, Part B3, pp. 165–175, 1996.