UPDATE OF ROADS IN GIS BY AUTOMATIC EXTRACTION FROM AERIAL IMAGERY*

C. Steger[1], W. Eckstein[1], C. Wiedemann[2]
[1]Forschungsgruppe Bildverstehen (FG BV), Informatik IX
[2]Lehrstuhl für Photogrammetrie und Fernerkundung
Technische Universität München
80290 München, Germany

ABSTRACT

Extracting roads from aerial imagery is important for the update of Geographic Information Systems (GIS), e.g., to capture new roads or to improve the geometric accuracy of existing roads. It can be broken down into two phases: First, old data is verified to determine which parts of the data have changed. Then, newly built roads are added. In this paper a two-phase approach consisting of automatic verification, as well as automatic road extraction is proposed. In the verification phase roads are searched for in the image along the axes given by the old data. This is done by using the gradient information of the image to localise possible roadsides. For the extraction of new roads a two-resolution approach with successive fusion has been developed. In a reduced version of the image center lines of roads are extracted, while in the original resolution roadsides are detected. The results of the two resolution levels are fused with a number of rules. Results of the proposed approach on aerial images are presented.

1 INTRODUCTION

This paper addresses the automatic update of road data in Geographic Information Systems (GIS). A procedure to do this task automatically can be broken down into two phases: First, old data needs to be verified to determine which parts of the data have changed (e.g., a road that has been removed due to a new bypass). Then, newly built roads are added to the data set.

A lot of work has been done in the area of automatic extraction of roads from aerial imagery (Airault, Ruskoné and Jamet 1994, Ruskoné, Airault and Jamet 1994, Barzohar and Cooper 1995) and the interactive, semi-automatic extraction of roads (McKeown Jr. and Denlinger 1988, Gruen and Li 1994, Heipke, Englisch, Speer, Stier and Kutka 1994, Vosselman and de Knecht 1995). While most of the automatic schemes try to extract roads by finding the corresponding roadsides in one way or another, many of the semi-automatic schemes use some sort of road profile tracing. None of the algorithms makes explicit use of á priori available GIS data, although many of the interactive systems could be modified to do so.

In comparison, there has been relatively little work in the area of verification of the old GIS data or its use to aid the extraction of roads (Moissinac, Maître and Bloch 1994, de Gunst and den Hartog 1994, Plietker 1994). Again, road profile matching and the extraction of roadsides are used to verify roads in the image.

*Presented at the Second International Airborne Remote Sensing Conference and Exhibition, San Francisco, California, 24–27 June 1996.

In this paper a novel algorithm to verify roads in high resolution aerial images (pixel size 20–50 cm) that uses roadsides as well as profile tracking is described. Furthermore, a two-resolution scheme to extract roads automatically from aerial images is described. It can be used to find the new roads after the verification.

2 VERIFICATION OF ROADS

In order to update GIS it is important to check which parts of the old data set have changed. One way is to compare the results of an automatic road finder with the old data. This has the disadvantage that the available data is not used to guide the road extraction. Furthermore, the extraction results have to be matched against the old data set, which is computationally very costly (Li, Kittler and Petrou 1992). Another way to check the old data which avoids this disadvantage is to look for the GIS roads directly in the image data. This has the following advantages: The area to be investigated is known; there is highly reliable information about the spatial position and, what is more, about the topology of the roads because most of the roads normally are unchanged. Using this information, it is possible to close gaps if they are enclosed by verified sections. By and large, there is a good chance to verify the roads using a simple model.

2.1 MODEL AND FUNDAMENTAL IDEA

The model comprises two fundamental assumptions about the appearance of roads in aerial images: Roads are assumed to have mostly straight and parallel roadsides. This means that if a road in the image corresponds to an axis of the old data, both roadsides will be approximately parallel to the axis. Furthermore, roadsides are assumed to correspond to strong edges in the image and the gray values along a road axis can be expected to be more or less constant. The fundamental idea of this approach is that both roadsides are expected to be near the axis if the old data corresponds to a road in the image. Thus, the first step consists in searching for the two strongest edges at both sides of the axis. This is done with loose constraints. For that reason some edges which are no roadsides will be detected. If the axis corresponds to the road the number of these false positions will be relatively small, otherwise many randomly distributed edges which are no roadsides will be found. The decision whether the axis corresponds to the road in the image is made in the second step using the following criterions: straightness and parallelism of the extracted edges and homogeneity of the gray values within the expected road.

2.2 VERIFICATION PROCEDURE

2.2.1 Edge Detection

To find the two strongest edges, the algorithm starts by computing a gradient image using a modified Deriche edge operator (Lanser and Eckstein 1992). Along each GIS road axis points with constant distance to each other are calculated. At these points relatively wide, symmetric profiles perpendicular to the axis are taken from the gradient image, similar to (McKeown Jr. and Denlinger 1988).

After this the positions of the two strongest edges within each profile are determined. The only constraint on the position of these two edge points within the profile is a minimum distance to each other. However, because of the loose constraints put on the edges, in general there will be a lot of outliers due to disturbances near to the road. This is also the case if the road axis in the GIS is inaccurate.

2.2.2 Width Estimation

Because of the outliers in the result of the edge detection it turnes out to be important to estimate the actual width of the road. This is done by calculating the center of the two edges and the distance of the center to the old axis for each profile. If this distance is less than a certain threshold (depending on the given level of accuracy), the two edge points are labeled as roadsides. The longest sections where the edge points are labeled as roadsides are calculated. For this task the imperfect sequence detector (ISD) described in (Aviad and Carnine Jr 1988) is used. For these sections the mean road width is estimated and used to adapt the width of the profiles. After adapting the width of the profiles to the road width, the search for the two strongest edges is repeated for each profile. By this means, disturbing edges further away from the road are eliminated. In Fig. 1 the result of the edge extraction process after the estimation of the road width is shown. For the case of relatively accurate GIS road axes (Fig. 1a) the algorithm is able to estimate the correct road width and to identify the true roadsides almost everywhere. In the case of inaccurate axes (Fig. 1b) the algorithm correctly estimates the road width, but cannot locate the correct roadsides for most of the inaccurate axis (see Fig. 2 for the axes).

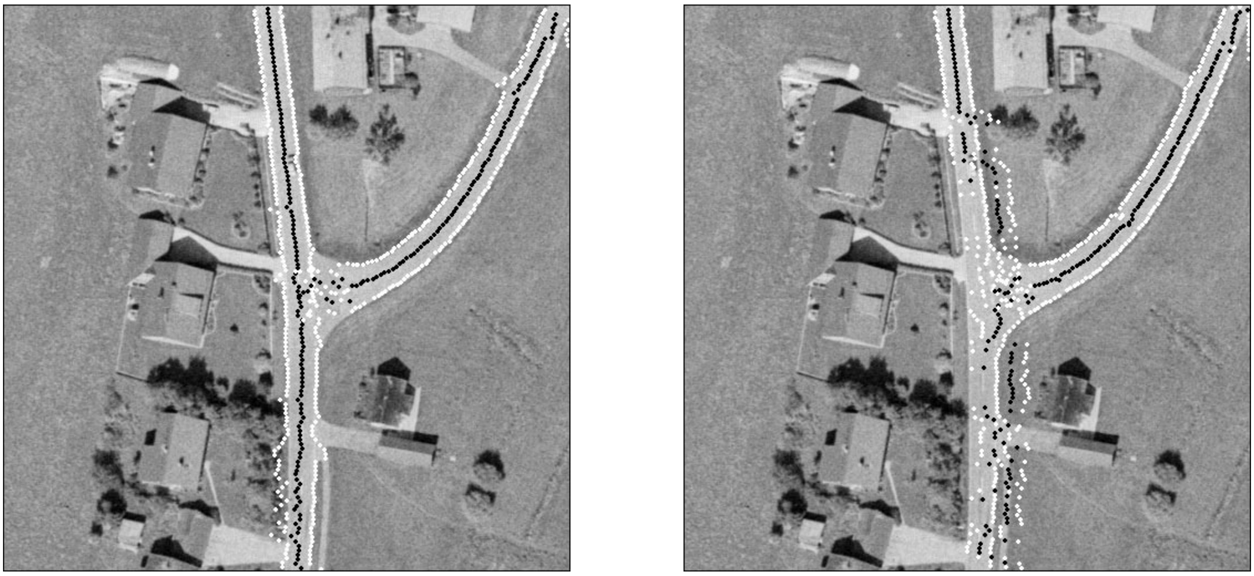

a) b)

Figure 1. Extracted edge points a) for accurate and b) inaccurate GIS road axes (see Fig. 2

2.2.3 Evaluation of the Old Axes

Two kinds of errors can occur in the labeling of the edge pairs: An error of the first kind is committed if the edge pair is labeled as not corresponding to the roadsides although both edge points correspond to them. An error of the second kind is committed if the edge pair is labeled as corresponding to the roadsides although this is not the case. These errors cannot be detected for each edge pair individually. It is done by checking the continuity of the extracted edge points along the direction of the axis.

A frequent reason for an error of the first kind is a slightly inaccurate position of the axis. This leads to a constant bias of the old axis and the center point of both edges. Therefore the edge pair will be labeled as not corresponding to the roadsides. This error typically occurs for many successive edge pairs. To detect this kind of error, the string of centers is checked for straightness along the road axis. Each point and its two neighboring points are connected consecutively by two vectors. The criterions for “straightness” are that the angle between the two vectors as well as the difference between the mean direction of the two vectors and the direction of the old axis are small. First all center points are labeled individually. Then it is checked if a gap in the string of edge points preliminary labeled as roadsides can be closed by a continuous string of center points labeled as straight. If this is the case, the corresponding edge points are labeled as roadsides as well.

The rate of errors of the second kind is calculated by checking all edge pairs which are labeled as roadsides. This is based on the criterions of straightness, parallelism and homogeneity. Typically roadsides are straight. Therefore all edge points which are colinear with their neighbors are assumed to be faultless and all others faulty. This is evaluated for each roadside separately. To check the edges for parallelism, the direction of the edge points is obtained from the direction image also calculated with the Deriche edge operator. A measure for parallelism of the two edge points within each profile is derived by comparing their directions.

It is not advisable to assume homogeneity of the gray values for the whole road as there are too many disturbances like cars or shadows. However, a great part of the road is homogeneous. What is more, an area depicting no road will often be distinguished by inhomogeneous gray values. The gray values of the center points of the road are assembled into a coarse histogram. The highest relative frequency will mostly be higher for roads than for other areas. Furthermore, the number of histogram sections with more than a certain threshold number of values will be less for roads.

Finally, all derived measures are integrated based on the product of all the calculated percentages to decide whether an old axis can be verified or has to be rejected.

2.2.4 Handling of Inaccuracies

Sometimes road axes are asymmetric or shifted with respect to the old axis in the image. Some parts of the axis lie within the road, whereas other parts do not. Typically there is a skip in the position of the detected edge points at the intersection of the old axis with the roadside. An example of this effect can be seen in Fig. 1b) in the upper center of the image. The edge which is intersected by the axis will be detected continuously. The corresponding roadside will only be detected if the axis lies between the two roadsides. There will also be a skip in the position of the calculated center points. To detect this skip, the gradients along the old axis are checked for significant high values. This results in a partitioning of the axis into several parts. To find out which part of the axis is lying between the roadsides, they are verified one after the other, using the algorithms explained in Sect. 2.2.3. Overall, this leads to a classification of the GIS roads into “verified,” “rejected,” and “inaccurate.”

2.2.5 Detection of Branching Roads

An important part of the verification task is to detect new roads. Each new road is connected with the existing road network. Hence there will be changes in the area around the old axes because of the new junctions. Often it is sufficient to check the immediate vicinity of the old axes to detect branching roads. For this task two different operators were developed. The first investigates the whole road to find new junctions whereas the second only checks areas when the roadside could not be recognized.

The first operator evaluates gray value profiles perpendicular to the old axis that are approximately symmetrical to the expected roadside. Normally the standard deviation σ will be high for each profile, as a part of it is lying outside the road. A sequence of profiles with low σ indicates a junction if its length is of the same magnitude of the expected road width.

The second operator uses the fact that often roadsides are not detectable in the area of junctions. In areas where the roadside could not be recognized a gray value profile is taken that is parallel to the old axis, but lying slightly outside the road. If this profile has a roadlike shape, i.e., if a bright region is detected, the center of this bright region will serve as the starting point of a profile perpendicular to the old axis. This profile connects the starting point with the old axis. If this new profile has more or less constant gray values a junction is expected.

2.3 RESULTS

Figure 2 shows the result of applying the algorithm to verify accurate and inaccurate road axes. If an axis is found to be correct it is displayed in white. Axes displayed in black are rejected. Dark gray axes are correct in principle, but are found to be geometrically inaccurate. As can be seen, all of the correct roads in the GIS have been verified correctly in Fig. 2a). In Fig. 2b) the inaccurate road is labeled as “inaccurate” as long as the axis lies within the actual road. Once the axis moves outside the road it is labeled as “rejected.” This corresponds with an intuitive notion about the correctness of a road axis. Furthermore, two possible new road starting points were found in the lower center and the bottom of the image in Fig. 2a), symbolized by short white lines. These examples show that the algorithm works well for different levels of accuracy of the GIS roads.

3 EXTRACTION OF ROADS

If roads were newly built or if the data set in the GIS is incomplete, the algorithm described in Sect. 2 will not detect any roads, although it will find some possible starting points of new roads. Therefore, this section will focus on the automatic extraction of roads from aerial images using a two-resolution scheme. Although the algorithm described here would benefit from using the starting point information that is obtained in the verification step, this has not been implemented so far.

3.1 EXTRACTION OF ROADS AT LOW RESOLUTION

At the low resolution level roads can mostly be modeled as lines that are brighter than their surroundings. Hence the first step to extract roads automatically is to extract lines. Because of the implicit smoothing performed in the scaling process, several kinds of problems that make the detection of roads at high resolutions more difficult are alleviated, e.g., cars on the road or shadows cast onto the road by adjacent trees. The algorithm can be described briefly as follows: The original orthophoto is reduced by a factor chosen such that roads in the reduced image are at most five to six pixels wide. From this image, lines are extracted by global thresholding, comparing the image to a local Gaussian mean, computation of connected contours, and selection of those contours that show a maximum perpendicular to their direction. A more detailed

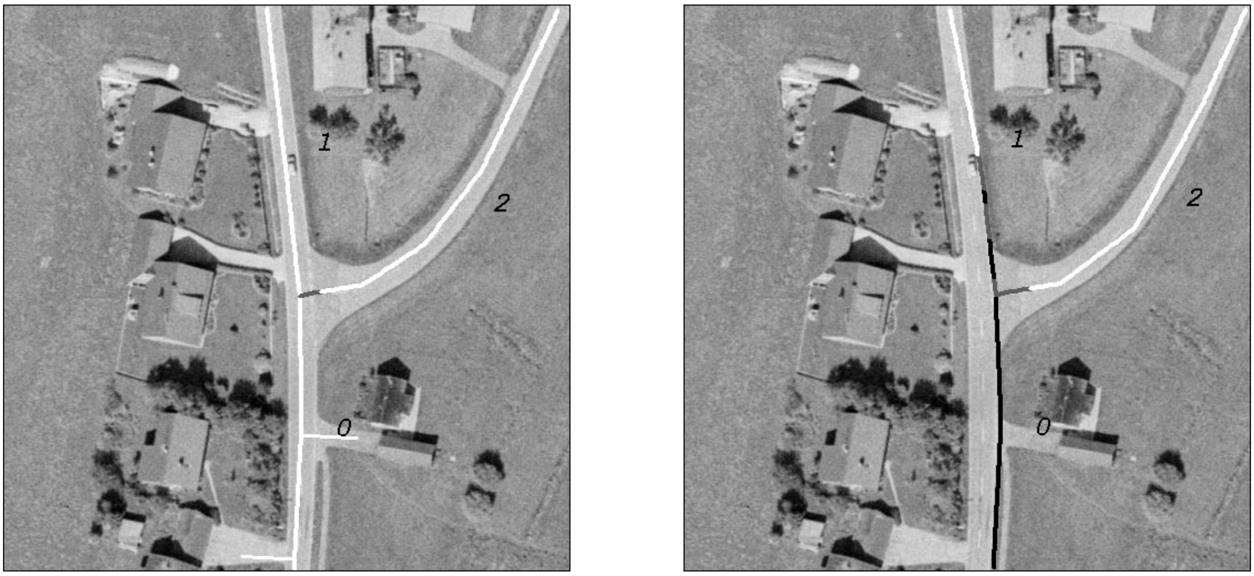

a) b)

Figure 2. Results of verifying a) accurate and b) inaccurate GIS road axes

description of the algorithm can be found in (Steger, Glock, Eckstein, Mayer and Radig 1995). The output of this step are the road hypotheses at the low resolution level. Figure 3 shows the result of this process.

3.2 EXTRACTION OF ROADS AT HIGH RESOLUTION

3.2.1 Edge-Extraction and Polygonal Approximation

At the high resolution level roads are assumed to be relatively homogeneous regions in the image that have a significantly different brightness than their surrounding areas. Hence it follows that roadsides can be detected by an edge extraction algorithm. As in Sect. 2.2.1, a modified version of a Deriche edge detector is used. The resulting edges are thinned by a non-maximum-suppression algorithm, yielding one pixel wide edges. From these edges contours are computed.

To reduce the amount of data which has to be handled, and to faciliate the perceptual grouping of parallel lines, a polygonal approximation of each contour is computed by the algorithm given in (Ramer 1972). This algorithm splits contours into polygon segments which have a limited distance to the approximated contour. The advantage of this algorithm is that it yields rather long line segments with an acceptably small approximation error.

3.2.2 Perceptual Grouping of Parallel Edges

One feature which characterises roads is the parallelism of opposite roadsides. Therefore the next step in the road-extraction process is to construct relations of parallel polygons. To be included into this relation, line segments have to fulfill several criteria. The first criterion is that the line segments should be approximately parallel. Parallelism in this case is defined by a threshold on the enclosed angle that depends

a)

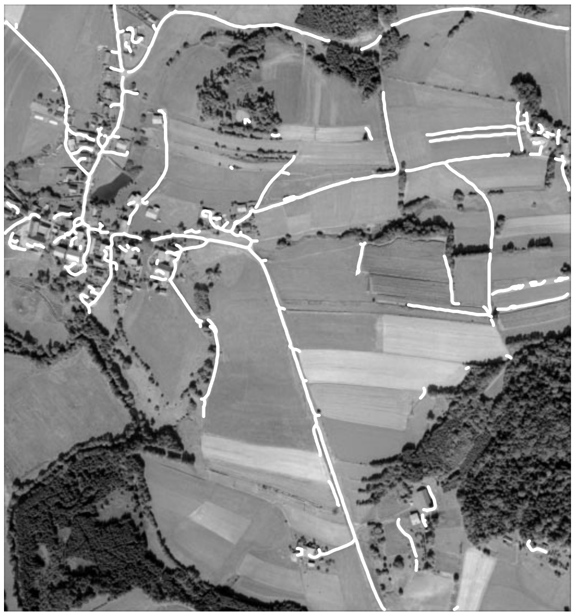

b)

Figure 3. a) Reduced version of the original orthophoto (reduced by a factor of 8) b) Result of the line detection algorithm

on the lengths of the two line segments involved. The threshold gets larger as the line segments get shorter. The reason for this is that roadsides in the image are never perfectly parallel and that longer line segments determine their direction more accurately. The second criterion requires parallel line segments to overlap. The last criterion is that the lines have to be closer than a certain threshold. This is motivated by the fact that roads cannot be wider than a certain distance. The exact details of the grouping procedure can be found in (Steger et al. 1995). Figure 4a) shows the result of this step.

3.2.3 Selection of Homogeneous Areas between Parallels

The next step in the algorithm is to examine whether the area between the parallel line segments found in Sect. 3.2.2 is homogeneous in the direction of the center line between them. This corresponds to the assumption that the surface intensity of a road is relatively constant in the direction of the road, whereas it can vary considerably across the road due to road markings or tire tracks.

To determine whether the region between two parallel line segments is homogeneous, slices parallel to the center line of the two lines are generated. These slices are 1 pixel apart, and the intensities within each slice are computed by bilinear interpolation. The mean intensity within each slice is computed. If the mean falls outside a certain user-settable range, the region is immediately rejected. This is due to the model assumption that a road only has gray values within a certain well-defined range for each image. After this the intensity variance in each slice is determined. This is the main criterion for homogeneity. If the variance in each slice is smaller than a threshold, the parallel lines that enclose the region are accepted as hypothetical roadsides.

The algorithm described so far fails in regions of the road where no parallelism can be detected, e.g., at intersections. Therefore all edges which neighbour the last parallel edge selected are examined if they still border homogeneous regions. To this end, rectangles are constructed that have the width of the last rectangle processed in the previous step. These rectangles are then sliced as above and examined for homogeneity. The result of this extension process are the road hypotheses at the high resolution level. An example can be seen in Fig. 4b).

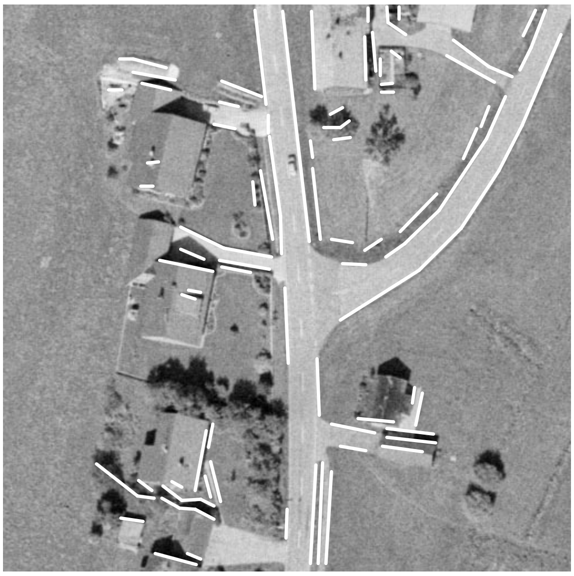

a)

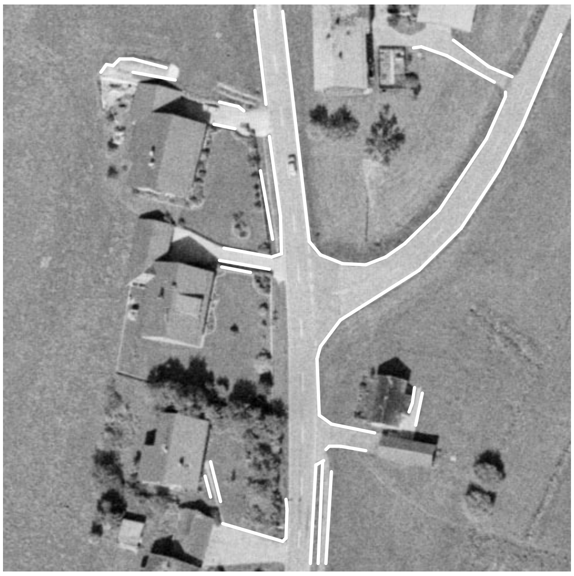

b)

Figure 4. a) Detected parallel line segments b) Selected parallels that enclose a homogeneous region or are adjacent to one

3.3 COMBINING THE DIFFERENT RESOLUTION LEVELS

As can be seen from the examples given so far, the extraction processes in the two different resolution levels have their advantages and deficiencies. Therefore in this section an algorithm is presented that combines the results of both levels to eliminate incorrectly detected road-segments and to extend results to regions where the extraction process has failed so far for some reason. The basic strategy in this step is to take the results of both levels which support each other. To start the process, parallel lines that enclose a homogeneous area of the high resolution level are selected if a center line of the reduced resolution level is found between the two parallels. Starting from these strong hypotheses, gradually all the roadsides are extracted. This is done using a number of rules. A detailed description of the rules is given in (Heipke, Steger and Multhammer 1995). The rules are rather conservative, selecting only edges as roadsides that have support in both resolution levels.

Figure 5a) depicts the input to the combination process. The results obtained in Sect. 3.2.3 are shown as solid lines, the original edges as dotted lines, and the center lines obtained in Sect. 3.1 as dashed lines. Figure 5b) shows the final result of the road extraction process. It can be seen that the algorithm was able to bridge the gap on the left hand side of the road at the intersection (see Fig. 4b) that resulted from lacking parallelism.

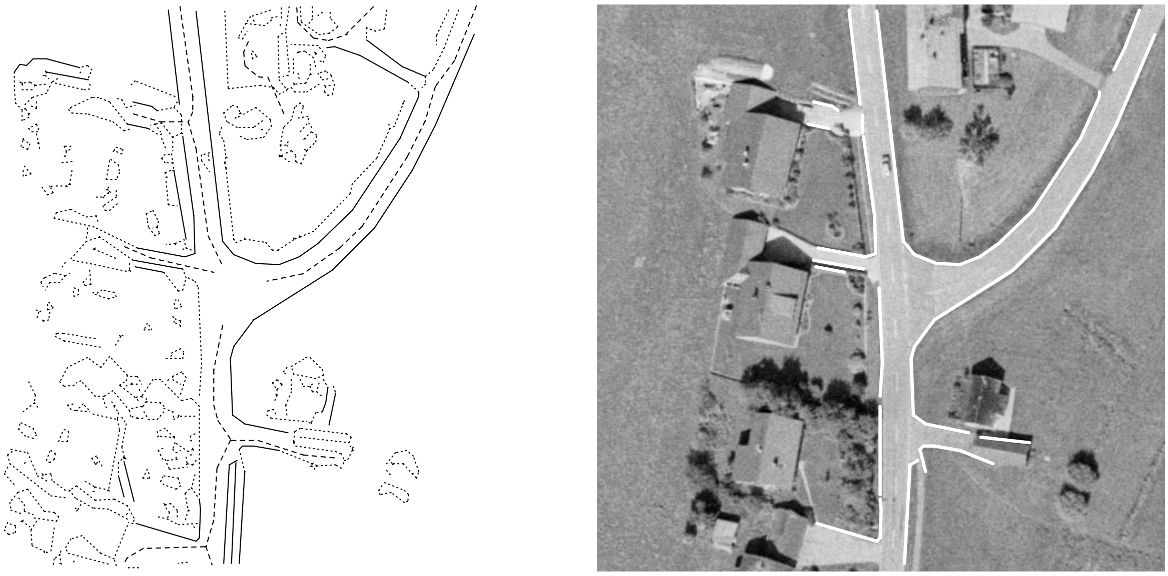

a) b)

Figure 5. a) Input to the combination step b) Final result of the road extraction process

4 CONCLUSIONS

Two algorithms have been presented that will be useful in implementing a system to automatically update roads in a GIS. The first algorithm is concerned with the verification of the old GIS data set. The experiments carried out with this algorithm indicate that it is able to classify the old road data correctly into “verified,” “rejected,” and “inaccurate” in most cases. Furthermore, it yields possible starting points for new road segments. The second algorithm extracts roads without prior data. It also gives acceptable results for a wide range of images. Further steps will focus on integrating these algorithms more tightly. The automatic road extraction algorithm could, for example, make use of the possible road starting points from the verification algorithm.

ACKNOWLEDGEMENTS

The funding of C. Steger by the German Science Foundation (DFG) under contract Eb 74/8–2 is gratefully acknowledged. We would like to thank H. Ebner, B. Radig, C. Heipke, H. Mayer, and A. Baumgartner for many valuable discussions.

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.

Aviad, Z. and Carnine Jr, P. D. (1988). Road finding for road-network extraction, *Computer Vision and Pattern Recognition*, pp. 814–819.

Barzohar, M. and Cooper, D. B. (1995). New geometric stocastic technology for finding and recognizing roads and their features in aerial images, *in* A. Gruen, O. Kuebler and P. Agouris (eds), *Automatic Extraction of Man-Made Objects from Aerial and Space Images*, Birkhäuser Verlag, pp. 255–264.

de Gunst, M. E. and den Hartog, J. E. (1994). Knowledge-based updating of maps by interpretation of aerial images, *12th International Conference on Pattern Recognition*, pp. 811–814.

Gruen, A. and Li, H. (1994). Semi-automatic road extraction by dynamic programming, *International Archives of Photogrammetry and Remote Sensing*, Vol. XXX, Part 3/1, pp. 324–332.

Heipke, C., Englisch, A., Speer, T., Stier, S. and Kutka, R. (1994). Semi-automatic extraction of roads from aerial images, *International Archives of Photogrammetry and Remote Sensing*, Vol. XXX, Part 3/1, pp. 353–360.

Heipke, C., Steger, C. and Multhammer, R. (1995). A hierarchical approach to automatic road extraction from aerial imagery, *in* D. M. McKeown Jr. and I. J. Dowman (eds), *Integrating Photogrammetric Techniques with Scene Analysis and Machine Vision II*, Proc. SPIE 2486, pp. 222–231.

Lanser, S. and Eckstein, W. (1992). A modification of deriche's approach to edge detection, *11th International Conference on Pattern Recognition*, Vol. III, pp. 633–637.

Li, S. Z., Kittler, J. and Petrou, M. (1992). Matching and recognition of road networks from aerial images, *in* G. Sandini (ed.), *Second European Conference on Computer Vision*, Vol. 588 of *Lecture Notes in Computer Science*, Springer-Verlag, pp. 857–861.

McKeown Jr., D. M. and Denlinger, J. L. (1988). Cooperative methods for road tracking in aerial imagery, *Computer Vision and Pattern Recognition*, pp. 662–672.

Moissinac, H., Maître, H. and Bloch, I. (1994). Urban aerial image understanding using symbolic data, *in* J. Desachy (ed.), *Image and Signal Processing for Remote Sensing*, Proc. SPIE 2315, pp. 310–321.

Plietker, B. (1994). Semiautomatic revision of street objects in ATKIS database DLM 25/1, *International Archives of Photogrammetry and Remote Sensing*, Vol. XXX, Part 4, pp. 311–317.

Ramer, U. (1972). An iterative procedure for the polygonal approximation of plane curves, *Computer Graphics and Image Processing* **1**: 244–256.

Ruskoné, R., Airault, S. and Jamet, O. (1994). Road network interpretation: A topological hypothesis driven system, *International Archives of Photogrammetry and Remote Sensing*, Vol. XXX, Part 3/2, pp. 711–717.

Steger, C., Glock, C., Eckstein, W., Mayer, H. and Radig, B. (1995). Model-based road extraction from images, *in* A. Gruen, O. Kuebler and P. Agouris (eds), *Automatic Extraction of Man-Made Objects from Aerial and Space Images*, Birkhäuser Verlag, pp. 275–284.

Vosselman, G. and de Knecht, J. (1995). Road tracing by profile matching and kalman filtering, *in* A. Gruen, O. Kuebler and P. Agouris (eds), *Automatic Extraction of Man-Made Objects from Aerial and Space Images*, Birkhäuser Verlag, pp. 265–274.