International Journal of Geographical Information Science
Vol. 22, No. 8, August 2008, 871–894

Research Article

Image-based quality assessment of road databases[X]

M. GERKE*† and C. HEIPKE‡
†ITC – International Institute for GeoInformation Science and Earth Observation, Hengelosestraat 99, P.O. Box 6, 7500 AA Enschede, The Netherlands
‡Institut für Photogrammetrie und GeoInformation, Leibniz Universität Hannover, Nienburger Str.1, 30167 Hannover, Germany

(*Received 13 March 2007; Final form 25 September 2007*)

In this paper an approach to the automatic quality assessment of existing geospatial data is presented. The necessary reference information is derived automatically from up-to-date digital remotely sensed images using image analysis methods. The focus is on the quality assessment of roads as these are among the most frequently changing objects in the landscape. In contrast to existing approaches for quality control of road data, the data to be assessed and the objects extracted from the images are modelled and processed together. A geometric-topologic relationship model for the roads and their surroundings is defined. Context objects such as rows of trees support the quality assessment of road vector data as they may explain gaps in road extraction. The extraction and explicit incorporation of these objects in the assessment of a given road database give stronger support for or against its correctness.

During the assessment existing relations between road objects from the database and extracted objects are compared to the modelled relations. The certainty measures of the objects are integrated into this comparison. Normally, more than one extracted object gives evidence for a road database object; therefore, a reasoning algorithm which combines evidence given by the extracted objects is used. If the majority of the total evidence argues for the database object and if a certain amount of this database object is covered by extracted objects, the database object is assumed to be correct, i.e. it is accepted, otherwise it is rejected. The procedure is embedded into a two-stage graph-based approach which exploits the connectivity of roads and results in a reduction of false alarms. The algorithms may be incorporated into a semi-automatic environment, where a human operator only checks those objects that have been rejected.

The experimental results confirm the importance of the employed advanced statistical modelling. The overall approach can reliably assess the roads from the given database, using road and context objects which have been automatically extracted from remotely sensed imagery. Sensitivity analysis shows that in most cases the chosen two-stage graph-approach reduces the number of false decisions. Approximately 66% of the road objects have been accepted by the developed approach in an extended test area, 1% has been accepted though incorrect.

*Corresponding author. Email: gerke@itc.nl
[X]Significantly enhanced and extended version of Gerke and Heipke (2006)

International Journal of Geographical Information Science
ISSN 1365-8816 print/ISSN 1362-3087 online © 2008 Taylor & Francis
http://www.tandf.co.uk/journals
DOI: 10.1080/13658810701703258

Those false decisions are mainly related to the lack of modelling road junction areas.

Keywords: Imagery; Networks; Quality; Reliability

1. Introduction

Geo-spatial data are the core and the most valuable part of any geographical information system (GIS). Information on its quality is important for both the data producer and the user. The quality control of such data generally can be seen to consist of three steps: check of logical consistency, verification and update. The logical consistency is ensured by comparing the data with the data model, i.e. the existence of attributes and the consistency of geometry etc. are checked. During verification and update the existing data are assessed using reference information, e.g. from remotely sensed imagery. In the verification step the geometric accuracy and the correctness of attributes (if available in the reference) are assessed. The completeness and temporal correctness are only partly considered, as only commission errors are identified. During a following update process, new or modified road objects not included in the database are extracted. By this means the completeness and temporal correctness are also fully considered.

The focus of our work lies on roads as these are among the most frequently changing objects in the landscape. In practical applications the update of road databases from imagery plays a minor role, since update information is often derived from other sources with a high update frequency. An example is the capture of roads using state-of-the-art mobile mapping vehicles or the direct import of planning data from the road construction administration. However, a periodic quality control of this data, i.e. verification, is necessary to avoid an accumulation of errors. This verification can be addressed very efficiently by using up-to-date orthoimages, either airborne or spaceborne. For this reason the verification using remotely sensed images is addressed in the current paper.

The background of this work is given by the WiPKA-QS-project carried out in conjunction with the German Federal Agency for Cartography and Geodesy (BKG). Here, a semi-automatic quality control system for the Authoritative Topographic Cartographic Information System (ATKIS) (AdV 1989, 2007) is developed. Object-based digital landscape models (DLMs) are important components of ATKIS. The ATKIS DLMBasis contains the data of the finest scale. Its content is approximately equivalent to topographic maps of scale 1:25,000, but in contrast to those maps the data contained in the DLMBasis are not subject to cartographic generalisation. Of course, a certain portion of generalisation is applied during data capture, but the important issue is that shapes and positions are in essence faithfully represented. Further information on the WiPKA-QS-project can be found in the work of Busch *et al.* (2004).

In our research only open landscape areas are considered, because these regions have two key advantages: (i) sophisticated and also practically relevant road extraction algorithms are available; and (ii) the modelling of relationships between roads and context objects, i.e. neighbouring objects, is of limited complexity and thus possible. The context objects, such as rows of trees, support the quality assessment of road vector data as they may explain gaps in road extraction. If for instance aerial images are captured in summer, trees along roads hamper the road extraction as the road surface is not directly visible. The extraction and explicit

incorporation of the context objects in the assessment of a given road database give stronger support for or against its correctness.

During the assessment, existing relations between road objects from the database and extracted objects are compared to modelled relations. If sufficient coincidence between model and extraction is observed, the database object is assumed to be correct, i.e. it is *accepted*, otherwise, it is rejected as *incorrect*. We do not verify the existing data in the classical sense of data maintenance where the reference has a clearly defined accuracy and reliability. The reference information in our case is obtained from automatic object extraction algorithms using remotely sensing images. Thus a portion of uncertainty is introduced – it is not clear beyond any doubt that the extracted objects are correct. The idea is to collect the information from these automatic procedures and compare it to a defined model. In this sense we compare two datasets and the actual quality of the data is not known. For this reason the intervention of a human operator still plays a key role during quality assessment. The automation should relieve the human operator in cases where the available information from the database and the extracted objects is not in conflict. A common and consistent modelling and processing of the road data to be assessed and the road objects extracted from the images is carried out. The term *quality assessment* is used instead of verification in the following to emphasise that the proposed approach aims at a statistical description of the quality of the given geospatial data.

The next section shortly introduces the related work, resulting in requirements for a new approach to road database assessment, which is described in section 3. Details on its realisation, results obtained with the new approach, a summary and an outlook are subject to the remaining sections of this paper.

2. Related work

This section gives a brief overview on work aiming at the verification of a given highly detailed road vector dataset. In the work of Gerke (2006) a comprehensive review on existing approaches to road network extraction and road database assessment is given; here only the core findings are summarised. Highly detailed road vector datasets are available for example in Germany (ATKIS DLMBasis), France (BDTopo), or Great Britain (OS Mastermap). Most of the existing approaches on road database verification (de Gunst 1996, Gerke *et al.* 2004) lack an adequate modelling of the relations between road and context objects, although it has been shown that the incorporation of context objects can give valuable evidence for road objects (Ruskoné 1996, Sardana *et al.* 2000). Moreover, the statistical properties of the input data are not sufficiently modelled and considered. Whereas in most approaches the uncertainty is not considered at all, the uncertainty is represented using a buffer of constant width in the work of Gerke *et al.* (2004). Buffering the objects implicitly assumes a uniform distribution for the geometric accuracy. This assumption however is not realistic: the normal distribution is more appropriate if only random errors affect object capture. Road network topology is incorporated in some approaches (Plietker 1997, Gerke et al. 2004) and was found to be a valuable means for enhancing the overall results.

These three aspects, namely the incorporation of context objects, a statistical modelling and the exploitation of network functionality, constitute requirements for an enhanced approach to road database assessment in open landscape areas.

Algorithms for the automatic extraction of road objects and rows of trees, representing the most salient context objects, from remotely sensed imagery in open landscape areas exist (Gimel'farb 1996, Straub 2003, Mayer *et al.* 2006). Therefore, adequate input information for the automatic quality assessment of road objects from geo-spatial databases is available.

3. Approach

The requirements which resulted from the discussion in the preceding section have been taken into account for the development of a new approach to road database assessment using remotely sensed imagery. This new approach consists mainly of four components (cf. figure 1). The *assessment algorithms* assess the GIS data through a quantitative comparison of the modelled relations (*relationship model*) and the existing relations between *extracted objects* and the given GIS road data. The *object model* provides a common geometric and statistical modelling of objects.

In the following, the modelling is first described in detail. Then, the object extraction is presented, and subsequently the overall workflow, which implements the strategy for road database assessment is outlined. The final subsection gives details on the developed algorithms for the object and network assessment.

3.1 *Models*

3.1.1 Relationship model. In the relationship model the geometric and topologic relations between a GIS road object, the local context objects and the extracted road objects are given (cf. figure 2).

The geometric relations *same shape* and *same orientation* express the fact that the course of a GIS road object and the respective extracted object needs to be

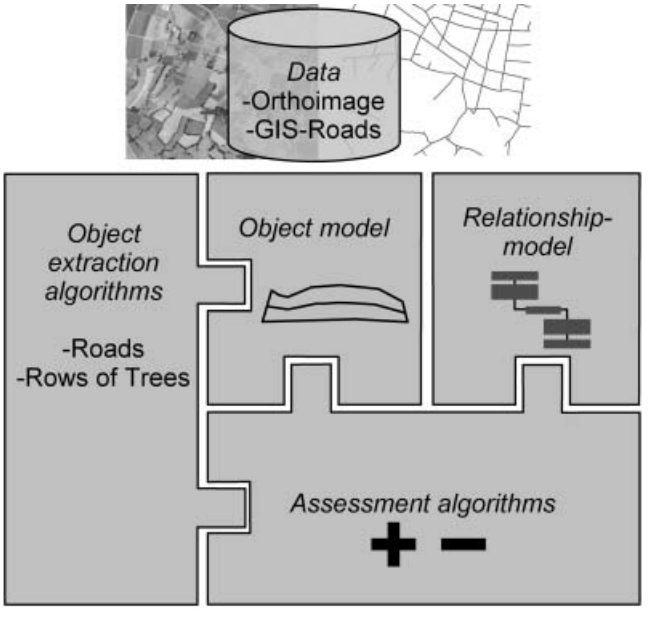

Figure 1. Components of the approach.

identical (shape) and that both objects must point towards the same direction. The topologic relation is important for this work due to the fact that, for example, rows of trees – more precisely the trunks of the trees – must be located outside the road given in the GIS database whereas an extracted road (the surface of the road) must be contained inside the GIS road surface. The topologic relations considered are *disjoint* and *contains*. The latter one is defined relative to the GIS object.

Besides these qualitative topologic relations we have defined a number of side conditions. For *disjoint* it is often desirable to define a minimum and a maximum distance ($d_{\min}$ and $d_{\max}$). For example, a row of trees must have a minimum distance to the road (due to security reasons) and it is also expected that trees having a distance to the road larger than a certain value are not relevant for explaining gaps in the road extraction, because they do not occlude the road in imagery. For *contains* a *same width* of objects may be required. The relations between the *GIS Road Object* and the *Local Context Object* and the respective values given in the depicted relationship model are defined by experience and common knowledge. An alternative way to find the measures would be to incorporate official specifications, for instance from road construction.

3.1.2 Object model – geometry. In the present approach roads are modelled as elongated objects. The model provides three geometric representations (refer also to figure 3):

(i) *pointset* Ⓢ;
(ii) *medial axis* M;
(iii) *border lines* $B_{L1,2}$, $B_{C1,2}$.Additionally, the attribute *width* is part of the object definition. Depending on the given data, it is assigned to the whole object as a constant or implicitly given as a variable distance measure between the border lines along the medial axis.

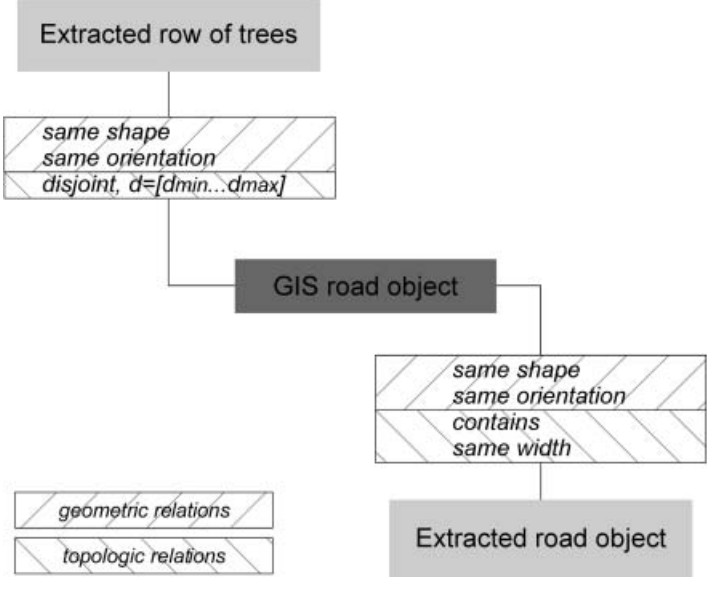

Figure 2. Relationship model.

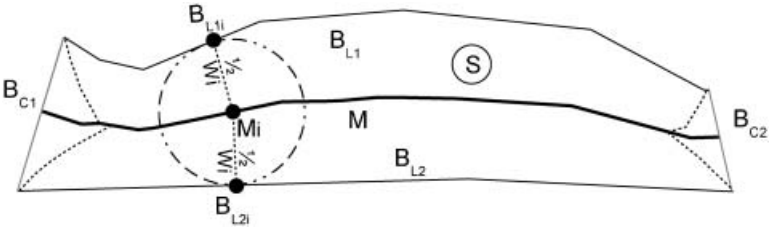

Figure 3. Object model: pointset Ⓢ, borders $B_{L1,2}$, $B_{C1,2}$ and medial axis M (continuous line), main branches of the skeleton according to definition in the literature: dashed line. The width W_i is assigned to M_i.

All representations may be defined both in discrete and Euclidean space. However, some definitions made here and subsequently used in statistical modelling are explicitly given in one of both spaces. Two representations are of particular importance for the assessment (refer again to figure 3): the pointset Ⓢ of the object is used to assess topologic relations and the medial axis M to assess shape and orientation. The border lines are necessary for the definition of the medial axis.

The skeletonisation of an object as described in literature (Blum 1967, Lee 1982, Rosenfeld 1986) has deficits for elongated objects. The skeleton does not completely cover the centre of the long-side borders (B_{L1} and B_{L2}): near the cross-sides (B_{C1} and B_{C2}) of such an object the skeleton is split (*see* dashed lines in figure 3[1]). However, for objects considered here, such a splitting is not desired. For a road object the medial axis should touch both cross-sides at their centre. Moreover, branches attached to the axis do not comply with the usual model of a road axis.

Therefore, we propose a modified definition for the medial axis. In the remainder the term *skeleton* refers to the definition presented in the literature and the term *medial axis* refers to the definition given in the following.

The medial axis M is defined in Euclidean space as follows: a point $M_i \in M$ is the centre of a circle with radius $\frac{W_i}{2}$ which touches the points $B_{L1i} \in B_{L1}$ and $B_{L2i} \in B_{L2}$, where B_{Ln} is the long-side borders of the object. The medial axis ends at the centres of the cross-side borders B_{C1} and B_{C2}. W_i is the width at M_i. The definition is similar to interpretations of a skeleton, for example the *grassfire process* (Soille 1999), but in this case the two fire fronts lit at B_{L1} and B_{L2} can only meet the opposite fire front, avoiding junctions and branches of the medial axis. Since the cross-side borders are not lit, the intersection of the fire fronts from B_{L1} and B_{L2} will finally propagate to B_{C1} and B_{C2}. For a detailed description of algorithms for the construction of M from the borderlines and the pointset refer to Gerke (2006).

3.1.3 Object model – statistics. In the statistical part of the model the object's positional uncertainty and imprecision are modelled through uniform and normal distributions.

Since remotely sensed images are used for object extraction, the first component in our statistical model is related to the necessary mapping of objects from object space to image space. We assume that the transformation parameters for this mapping are known accurately enough. In practice this means that we are working in imagery, where a height model of the terrain was used to orthorectify a given aerial or satellite imagery. These orthoimages constitute standard products in

[1] In order to avoid too many details in the figure, additional branches of the original skeleton are not shown.

remote sensing data processing. In so-called *true orthoimages* the rectification was done incorporating a digital surface model, i.e. also the height of objects above the terrain. However, the extraction of rows of trees, being 3D-objects, from orthoimages where only a digital terrain model was considered for the orthorectification leads to an offset in the x-y-plane. For our approach the resulting translation of the object's position in 2D has to be considered in the statistical modelling. In order to simplify the modelling, the 3D-object is assumed to have a constant height. For the case of aerial images and central perspective projection equation (1) describes the offset (Kraus 1993).

$$\Delta R = \frac{H}{c/\rho + \tan\alpha}, \tag{1}$$

where ΔR stands for offset in the x-y-plane in the object coordinate system, H for height of the object, c for calibrated focal length, ρ for distance of object from the nadir point in image space and α for ground slope at the position of the object. In our model ΔR is introduced as the radius of a uniform distribution in case no further knowledge on the position of a 3D-object with respect to the image centre is available.

Besides the uncertainty caused by the described unknown mapping from object space to image space, the uncertainty of object geometry is related to the fuzziness of abstraction (Winter 1996, 1998). An example is the decision of an operator concerning the question which points on the road surface belonging to the road's medial axis. The imprecision is related to the measuring process, i.e. it is the deviation of a (unknown) true variable from its mean (measured) value. The individual parameters for the probability density functions, i.e. standard deviation for the normal case and radius for the uniform case, need to be chosen according to the used input data for object extraction and specific object properties. The final density function of the object is obtained by a convolution of the input functions.

In table 1 the statistical components of the model are listed. For each component the assumed probability density function (PDF) is given in the following notation: D_X is the radius of a uniform distribution and $G_{\sigma X}$ the standard deviation of a Gaussian distribution.

The fourth and fifth columns indicate whether the respective component has an impact on the geometric and/or the topologic relation between objects. One can see that two components do not have an influence on the geometric relations: *Object modelling* and *Mapping object → image*. The geometric relations are not influenced,

Table 1. Statistical components of the object model.

	PDF and its parameter		Impact on …relationship		Concerns …objects	
Component	Normal	Uniform	Geometric	Topologic	GIS	Extracted
Object modelling		D_{Mo}	No	Yes	Yes	No
Mapping object → image		D_O	No	Yes	No	Yes
Abstraction	$G_{\sigma A}$	D_A	Yes	Yes	No	Yes
Measurement	$G_{\sigma C}$		Yes	Yes	No	Yes

because the assumed constant position offset does not modify the shape and the orientation of the object.

The last two columns indicate which component has an impact on which type of object, i.e. on the road database objects or on the extracted objects. The modelling only has an impact on road database objects. The influence of the other components cannot be modelled explicitly as generally no quantitative information on the data capture process is available. On the other hand, the extracted objects are assumed to be captured from imagery and thus all related components influence the statistical modelling. The object modelling is not considered here, as it is normally impossible to quantify it. However, it is implicitly contained in the other three components.

The actual values for the statistical parameters depend on many factors and cannot be given in general. For example, if rows of trees are extracted from an orthorectified aerial image, where the original exterior and interior orientation parameters are known and the tree heights can be reasonably estimated, the offset can be computed (cf. equation (1)) and thus the planimetric position can be corrected. Then D_O is approximately zero and thus the offset can be neglected. However, if such information is not available, D_O needs to be estimated based on the available knowledge of image acquisition and tree type.

To simplify the statistical modelling, the following assumptions are made:

(i) uncertainty is assigned to every point $P_i \in \mathbb{S}$;
(ii) uncertainty is assumed to be identical for every point $P_i \in \mathbb{S}$ and in both coordinate directions;
(iii) points of the borders are a subset of $\mathbb{S}$ and therefore have the same uncertainty as every other P_i.

These assumptions are reasonable as the objects considered are captured from digital imagery. If information on systematic effects is available—for instance, for distortions of the optical system—these are normally corrected in the processing and thus do not need to be considered in our approach.

3.2 *Object extraction*

Roads are extracted using the approach described by Wiedemann (2002). They are modelled as linear objects in aerial or satellite imagery with a resolution of about 1–2 m. The underlying line extractor is the one introduced by Steger (1998). The initially extracted lines are evaluated by fuzzy values according to attributes, such as length, straightness and constancy in width and grey value. The final step is the grouping of the individual lines to derive topologically connected and geometrically optimal paths between seed points. The decision of whether extracted and evaluated lines are grouped into one road object is made according to a collinearity criterion, allowing for a maximum gap length and a maximum direction difference.

The rows of trees are automatically extracted as follows: the supervised multi-scale, multi-spectral segmentation approach of Gimel'farb (1996) is applied for the segmentation of the scene into the classes built-up, industrial, grassland, cropland and forest. The required training data are defined by a human operator. Within the objects of the forest class rows of trees are chosen as elongated regions. Regions having an average width larger than the average crown diameter as manually estimated for the scene are not further considered. The medial axes of the objects are derived according to the model outlined above. In figure 4 an example is given, showing the extracted objects.

The employed procedure assumes implicitly that the medial axis of a row of trees segment is a good representation for the line connecting the trunks. Here, the fuzziness of abstraction becomes important: whether or not the trunk is centred beneath the crown can normally not be observed in aerial imagery. Therefore, the actual values for the distribution describing the fuzziness of abstraction need to be chosen adequately.

4. Workflow and assessment algorithms

The process of road database assessment consists of three main parts: *Initial object assessment*, *Graph analysis* and *Final object assessment*, refer to figure 5. This partitioning makes the network exploitation aspect explicit: the *Initial* and the *Final object assessment* are related to individual road objects of the GIS database, whereas in the graph analysis the whole network is considered. In the following text, the term *Phase 1* is used for the *Initial object assessment*, while the term *Phase 2* is employed for the *Final object assessment*. Phase 1 aims at a very reliable assessment. This means that the number of falsely accepted road objects (false positives) should be very small. As a consequence, the total number of accepted road objects is also quite small. In the graph analysis part the network function is exploited, and local context objects are also incorporated. The idea is to identify those roads which have been rejected in Phase 1 but which perform important network functions, i.e. connect reliably extracted road network components. If local context objects give hints regarding the correctness of those important road objects, they will be accepted, otherwise, they are subject to a second assessment in Phase 2. In that phase road extraction is applied a second time, but the algorithm uses a more tolerant set of parameters for road extraction.

The following assessment algorithms are applied to individual GIS road objects. In Phase 1 every road object stored in the database is assessed, whereas in Phase 2 only those objects are regarded which turned out to fulfill an important network function, but which have been rejected in Phase 1, and for which local context objects could not explain the rejection.

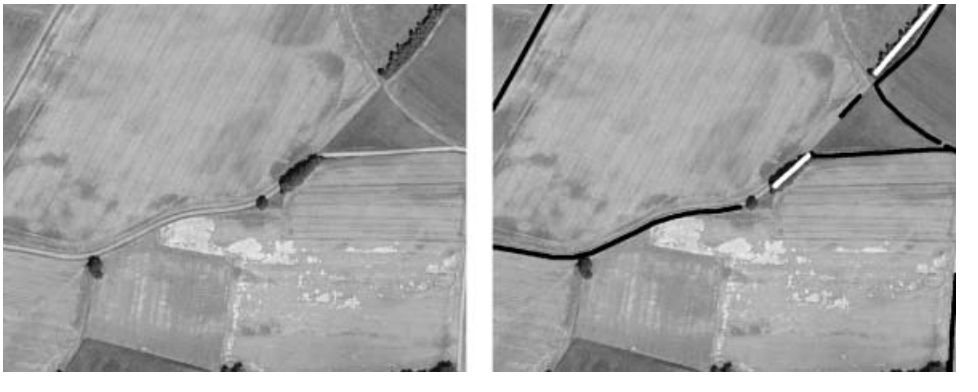

Figure 4. Object extraction. Left: image, right: extracted roads (black) and extracted rows of trees (white).

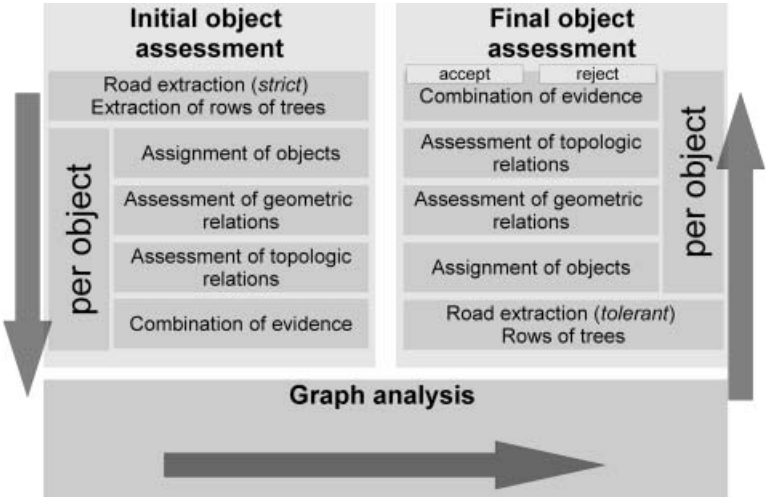

Figure 5. Workflow.

4.1 *Assignment of objects*

The task here is to find those extracted road objects and local context objects (or parts thereof), which may give evidence whether the currently processed GIS road object maintains the relations modelled in the relationship model. Additionally, the coverage $qcov_i$ between the GIS object and the extracted object is calculated. This value is used as a weight factor indicating the influence an extracted object has on the overall assessment of the GIS object.

4.2 *Assessment of geometric relations*

In this component a probability P_{g_i} is calculated. P_{g_i} expresses to what extent an extracted object i and the considered GIS road object maintain the modelled geometric relations *same shape* and *same orientation*. The probability P_{g_i} consists of two components: the extent to which the shapes of both objects are identical is expressed by $P_{g-shape_i}$, while P_{g-ori_i} concerns the identity of the orientation. Both values are regarded as independent. Thus, the final value is the product of both measures:

$$P_{g_i} = P_{g-shape_i} \cdot P_{g-ori_i}. \tag{2}$$

The medial axis representation is used for the calculation of both components.

4.2.1 Assessment of shape identity. The basic idea of comparing the shape of two objects is that if the shapes are identical, the translation and rotation invariant line moments need to be identical, too. This choice is based on the moment uniqueness theorem (Hu 1962) which states that an object can be represented by an (infinite) set of moments. Details on line moments are given in appendix A.

When the uncertainties of the objects are considered, the comparison of the moments can be expressed in terms of a statistical test. The variance propagation which is necessary to transfer the probability measures from the line object to the moments was developed by Gerke (2006). The probability P_{pq} that the translation and rotation invariant moments μ_{pq} of two objects are equal is expressed by the

probability that the difference of both moments is zero:

$$P_{pq} = F\left(y_{1-\alpha/2} - \widetilde{D}_{pq}\right) - F\left(y_{\alpha/2} - \widetilde{D}_{pq}\right), \tag{3}$$

where $y_{1-\alpha/2}$ and $y_{\alpha/2}$: boundaries of the confidence interval (quantiles of the standard normal distribution). Here the probability of error $\alpha=0.01$ is chosen. $\widetilde{D}_{pq}$ is the test statistic, i.e. the difference of moments divided by its standard deviation. F is the normal distribution function. The fact that P_{pq} equals the area under the density function of $\widetilde{D}_{pq}$ between the confidence boundaries is illustrated in figure 6.

Finally, the probability $P_{g-shape_i}$, considering all moments with order o from o_{min} to o_{max} is

$$P_{g-shape_i} = \coprod_{o=o_{min}}^{o_{max}} \coprod_{q=0}^{o} P_{pq} \text{ with } p = o - q. \tag{4}$$

The moments of the first and second orders are used for the calculation of the centre of gravity and the orientation. For this reason, the minimum order used for the comparison of translation and rotation invariant moments is $o_{min}=3$. In empirical studies it was found that moments of order larger than 8 are not significantly different from zero in real data, therefore, we set $o_{max}=8$.

4.2.2 Assessment of orientation identity. P_{g-ori_i}, the probability that both orientations are identical, is computed applying a similar approach as developed for the determination of P_{pq}. Here, the test statistic is derived from the difference of the orientation of both objects. The object's orientation and the respective variance are obtained from line moments, again see Gerke (2006) for more details.

Example for the assessment of geometric relations.

The GIS road object and the extracted road object shown in the left hand image of figure 7 have a similar shape. The right side of figure 7 shows the probabilities $P_{g-shape}$ and P_{g-ori} as a function of the standard deviation for the imprecision of measurement G_{σ_C}. The value for G_{σ_C} has been varied from 0.01 to 1 m, thus a logarithmic scale is used for the right hand side of the x-axis of figure 7.

The identity of both shapes is reflected in a large value for $P_{g-shape}$: for $G_{\sigma_C}=0.01$ m the value is zero, it increases to 0.4 for $G_{\sigma_C}=0.02$ m and to 0.9 for $G_{\sigma_C}=0.05$ m, and finally for $G_{\sigma_C}\geq 0.1$ m the probability that both shapes are identical is 0.99. According to the theory of line moments this is an expected result, because both objects are largely straight and parallel, i.e. deviations in the shape are restricted to some short segments which have no significant impact on the

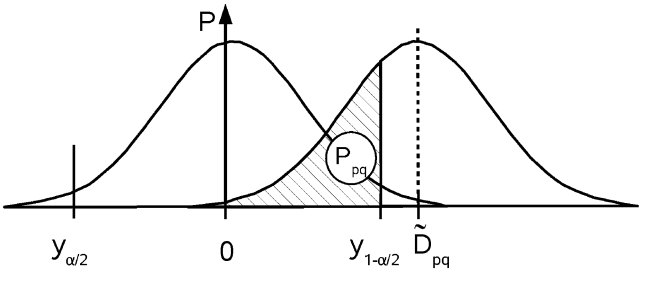

Figure 6. Calculation of P_{pq}.

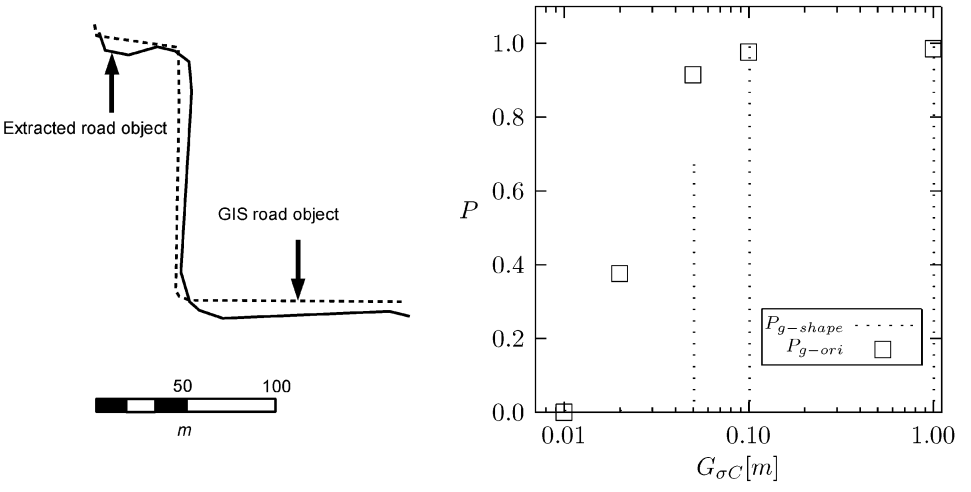

Figure 7. Example for assessment of geometric relations: calculation of $P_{g-shape}$ and P_{g-ori} for a varying assumed measurement accuracy.

differences of line moments. For P_{g-ori} a similar observation can be made: for the highest assumed accuracy of the given data, the probability that both objects point towards the same direction is zero, but with increasing G_{σ_C} the probability that both lines have the same orientation also increases; it is 0.99 for $G_{\sigma_C} \geq 0.1$ m. This is also reasonable as the deviations in the pointing direction between both lines are small compared to the overall length.

This example draws some conclusions: the developed approach for the assessment of geometric relations reasonably reflects the shape and orientation identity of two objects under the consideration of imprecision. However, in certain circumstances the result from the analysis of shapes based on the proposed method may differ from subjective observations of a human operator. Nevertheless, because of lacking alternatives we use reference data which were prepared by a human operator for the test of the overall approach (cf. section 5).

4.3 *Assessment of topologic relations*

Topologic relations describe the relative position of two objects, not considering their metric; they are invariant with respect to many geometric transformations. A significant contribution to this field has been given by Egenhofer and Franzosa (1991). For two objects (regions) in R^2, being singularly connected and regularly closed, a set of eight topologic relationships can be distinguished: *disjoint*, *touch*, *overlap*, *contains/contained by*, *covers/covered by*, *equal*.

The task in this step of our approach is to determine the probability P_{t_i} that the given GIS road object and an extracted object keep the modelled topologic relation. For the examination of the topologic relations the approach presented by Winter (1996, 1998) is applied and extended. In that work the topologic relations between imprecise and uncertain regions are assessed, considering an arbitrary probability density function for the respective object's borders. Winter shows that all eight topologic relations can be derived from the minimum and maximum distance between so-called certain zones of both objects. The relations modelled in our

approach – *contains* and *disjoint* – are assessed using this method. Again, details are given by Gerke (2006).

4.4 *Combination of evidence*

The objective of the combination of evidence is to find a quality indication for the GIS road object incorporating the results from the assessment algorithms just described. Every extracted object, i.e. roads or rows of trees, being assigned to a GIS road object allows a conclusion $\xi_i=1$, $i=1, \ldots, N$, $N=$the number of objects being assigned to a particular GIS road object, which states that the GIS object and the respective extracted object keep the modelled geometric and topologic relations. The probability of whether $\xi_i=1$ is true $\left(P_i^+\right)$ or false $\left(P_i^-\right)$ is assumed to depend on the collected measures.

The probability P_{g_i} that the modelled geometric relation holds gives the main evidence: if it is likely that the shape and the orientation of both objects comply with the model, P_i^+ will be correspondingly large, and vice versa: the smaller P_{g_i}, the larger P_i^- will be.

The measures P_{t_i} and $qcov_i$, i.e. the amount of coverage between extracted object and GIS road object, describe the impact an extracted object has on the assessment of the respective GIS road object: the larger these values are, the larger the evidence given by P_{g_i} should be. Thus, these two measures are interpreted as weight factors α_i. Under the additional assumption that no other influences exist these considerations lead to:

$$\alpha_i = P_{t_i} \cdot qcov_i, \tag{5}$$

$$P_i^+ = P_{g_i} \cdot \alpha_i, \tag{6}$$

$$P_i^- = \left(1 - P_{g_i}\right) \cdot \alpha_i. \tag{7}$$

The incorporation of P_{g_i} and $\left(1-P_{g_i}\right)$ in both P_i^+ and P_i^- implies a threshold of 0.5: if $P_{g_i} > 0.5$, the evidence given by the object i is in favour of P^+, otherwise, P^- is larger than P^+.

In order to be able to assess the whole GIS road object, the evidence delivered by all extracted objects assigned to the GIS road object need to be combined. Two hypotheses θ_1 and θ_2 are defined for this purpose:

$\theta_1 = H^+$: the whole GIS road object is correct given the observed data, i.e. the modelled relations hold for the extracted objects and the GIS road object, and

$\theta_2 = H^-$: the whole GIS road object is not correct given the observed data, i.e. the modelled relations do not hold for the extracted objects and the GIS road object.

An approach combining all conclusions $\xi_1 \ldots \xi_N$ related to, i.e. giving evidence for, a GIS road object must consider all the individual probabilities and finally infer the quality, permitting an overall assessment conclusion, i.e. confirm H^+ or H^-.

The conditional probabilities for the correctness of the statement $\xi_i=1$ are given by P_i^+ and P_i^- and the weight factor α:

$$P(\xi_i = 1 | \theta_1 = H^+) = P_i^+ = P_{g_i} \cdot \alpha_i, \tag{8}$$

$$P(\xi_i = 1 | \theta_2 = H^-) = P_i^- = \left(1 - P_{g_i}\right) \cdot \alpha_i. \tag{9}$$

The ξ_i is assumed to be independent:

$$P(\xi_i, \xi_k|\theta_j) = 0 \qquad \forall i \neq k, j = [1, 2], \tag{10}$$

therefore, the combined probability for the correctness of θ_1 and θ_2 can be derived by adding the individual probabilities:

$$P(\xi_1 + \ldots + \xi_N|\theta_j) = \sum_{i=1}^{N} P(\xi_i|\theta_j). \tag{11}$$

The assumption of independence is a quite optimistic one, because it may be expected that a certain correlation exists between objects extracted by the same algorithm from the same image. However, as every extracted object gives evidence for H^+ and H^-, the influence of possible correlations on the probability for both hypotheses may be assumed to be similar, and thus has no significant impact on the final maximum probability decision.

Whether the road database object is accepted depends on a maximum probability decision. Additionally, a required minimum total coverage needs to be reached for the road database object to be confirmed. This is important to assure that a major part of the GIS object was assessed.

In the overall strategy, the separation of evidence from extracted road objects and local context objects is a key issue. In Phase 1 an object is *fully accepted* if solely the extracted road objects confirm its correctness; it is *preliminarily accepted* if only the fusion of evidence from extracted road objects and local context objects lead to a confirmation of the hypothesis H^+. This decision implies that the local context objects are necessary to explain gaps in the road extraction. Finally, if none of the extracted objects gives enough evidence to confirm H^+, the respective object from the database is *preliminarily rejected*.

In Phase 2 only the *preliminarily rejected* road objects performing important network functions, i.e. connecting reliably extracted road network components, are checked again, leading to a *finally accepted* or *finally rejected* decision, respectively. Whether a particular GIS road object performs important network functions is decided in the graph analysis step, which will be described in the next section.

In figure 8 the possible transitions concerning the assessment results are shown.

4.5 *Graph analysis*

The steps to road object assessment described in the preceding sections are related to individual road objects; the road network has not been incorporated so far. The exploitation of road network topology is based on the assumption that road objects having been accepted in Phase 1 need to be connected and form a network. This is a reasonable assumption, because isolated roads occur unlikely in reality. To restrict the number of possible connections and according to the assumption that road objects link places on short paths, the criterion to be fulfilled is that the distance between objects is minimised. The further treatment of objects not having been *fully accepted* in Phase 1 depends on whether they are part of the shortest distance between *fully accepted* objects and on the amount of evidence given by local context objects in Phase 1. Note that although it is not certain whether the network given in the database is fully correct, it is used to formulate such connection hypotheses here.

The assessment result per GIS road object is transferred to edges of the road network graph. This is important to be able to reasonably analyse the network

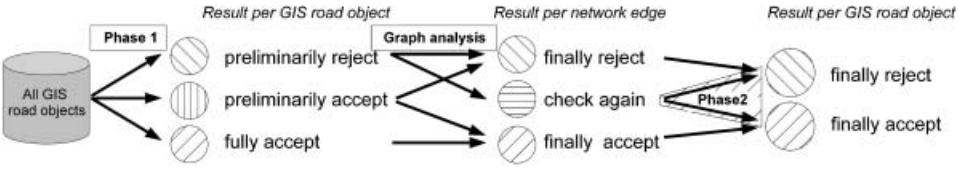

Figure 8. Possible transitions between assessment results concerning the three parts: initail object assessment (Phase 1), graph analysis and final object assessment (Phase 2).

function of the assessed roads. Afterwards, the shortest paths between the *fully accepted* edges are derived, applying the A*-Algorithm (Duda and Hart 1973). The edges are then labelled according to the following rules: if a *preliminarily accepted* edge is part of a shortest path, it is labelled as *finally accept*. This rule is motivated by the assumption that the local context objects are able to explain existing gaps in the road extraction. This assumption is additionally supported by the network analysis and thus the correctness of the edge – respectively of the assigned GIS objects – may be expected. If a *preliminarily rejected* edge is part of a shortest path, it is labelled as *check again*. The network function of those edges is assumed to give enough evidence to check them again in Phase 2, applying a road extraction with more tolerant parameters. All remaining *preliminarily rejected or accepted edges* are labelled as *finally reject* as it is assumed that they do not fulfill an important network function. The background of this decision is to minimise the false positives, therefore, it is demanded that support from the network function must be given to check objects in Phase 2.

The labels from the edges are then transferred back to the GIS road objects. The objects being labelled as *check again* are prepared to be processed again in Phase 2, the remaining objects, i.e. objects being labelled *finally accept* and *finally reject*, are not further considered.

5. Test of the approach

The described approach was implemented and tested. The chosen test site for this paper is located in the German state North Rhine-Westphalia. It has a size of $2 \times 8\,\mathrm{km}^2$. The employed RGB orthoimages have a ground sampling distance of about 30 cm. The ATKIS DLMBasis contains about 530 road objects. A manual reference verification of the dataset showed that about 98% of the objects are correct, and about 25% of them are significantly occluded by rows of trees. The object extraction from the given imagery was conducted using the algorithms described in section 3.2. Figure 9 gives an impression of the scene. Individual road objects cannot be identified in detail in the image. However, the spatial arrangement of the objects is very well visible.

A means to evaluate the assessment results is to define a confusion matrix, where the reference and the assessment result are compared; refer to table 2. In this matrix the columns indicate the result from the automatic assessment approach, i.e. *accept/reject* and the rows show the manual reference decision, i.e. *correct/incorrect*. The types of error and their impact on the practical semi-automatic process can easily be identified using this matrix. The operator who has to inspect the road assessment results needs only to concentrate on the objects which have been rejected, i.e. those in the right column. Therefore, the number of true positives should be large since it indicates efficiency. The false positives are undetected errors and thus should be minimised.

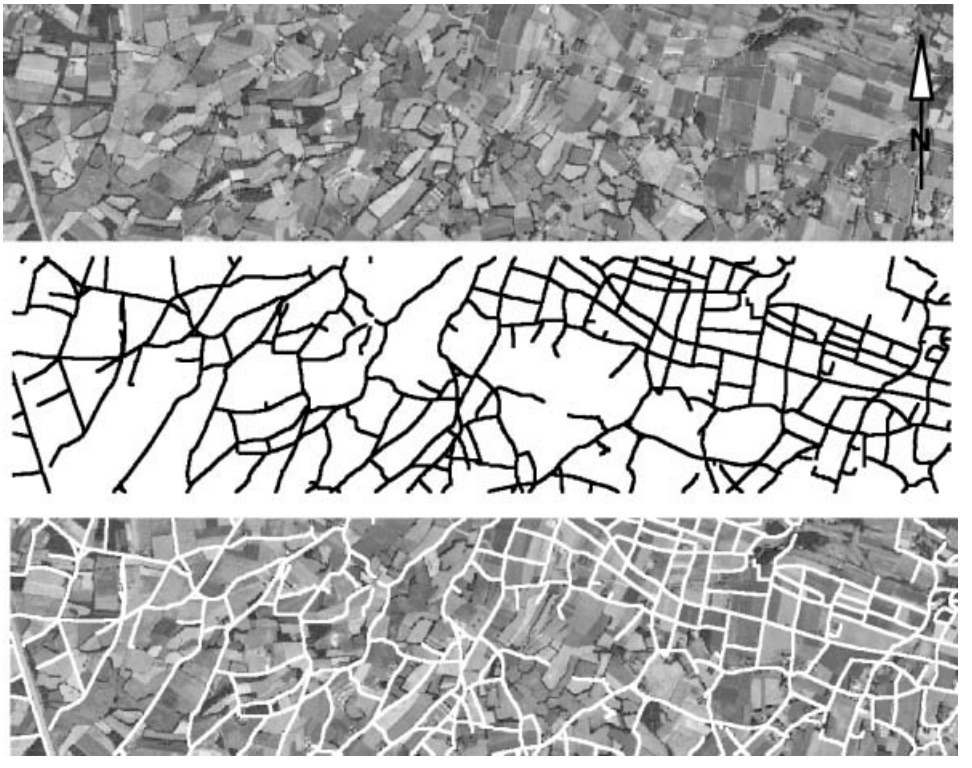

Figure 9. Top image: orthoimage, centre image: road objects from ATKIS DLMBasis, bottom image: ATKIS road objects superimposed to orthoimage. Orthoimages and ATKIS DLMBasis © VA Nordrhein-Westfalen, Germany, 2006.

5.1 *Assessment of road objects – without incorporating the rows of trees*

First, the network approach was applied, but no rows of trees were incorporated. The confusion matrices are shown in table 3. The left confusion matrix is related to all objects. It shows that nearly 60% of the road objects are correct and have been accepted, while about 38% have been rejected although being correct. All together 2.1% of the objects have been labelled as incorrect by the operator, and approximately half of those errors were identified automatically. A detailed discussion of false positives and false negatives is given below.

To emphasise the influence of the trees regarding the assessment, the confusion matrix on the right side only relates to road objects which are significantly occluded

Table 2. Confusion matrix for the evaluation of results.

		Assessment decision	
		Accept	*Reject*
Reference	*Correct*	True positive (efficiency)	False negative (manual postprocessing)
	Incorrect	False positive (undetected error)	True negative (manual postprocessing)

Table 3. Confusion matrices: assessment result without the incorporation of rows of trees.

		Assessment result	
		Acceptance	*Rejection*
Reference	*Correct*	59.8%	38.1%
	Incorrect	1.0%	1.1%

Consideration of all objects

		Assessment result	
		Acceptance	*Rejection*
Reference	*Correct*	52.8%	47.2%
	Incorrect	0.0%	0.0%

Consideration of only the objects being significantly concluded by trees

by trees. Only approximately 53% of true positives are obtained, compared to approximately 60% shown in the left matrix, i.e. the occlusion leads to an insufficient road extraction in a number of cases. A further observation is that none of those road objects occluded by rows of trees had been labelled incorrect by the human operator, and please see the lower line in the right hand matrix.

5.2 *Assessment of road objects – incorporating the rows of trees*

In the second test the complete workflow as proposed in this paper is utilised. During Phase 1 extracted roads are used to preliminarily assess the given ATKIS road network. The automatically captured rows of trees are also incorporated, i.e. if a road is rejected, but connects two important network nodes, it will be accepted as far as rows of trees give enough evidence; otherwise, it will be selected for the final object assessment (Phase 2). In Phase 2 the selected roads are assessed again, but using a more tolerant parameter set for road extraction.

The confusion matrices are shown in table 4, again related to all objects (left) and related only to significantly occluded road objects (right).

The number of true positive decisions increases from nearly 60% (network exploitation, without rows of trees, cf. table 3) to 66%. Furthermore, one additional object has been classified as false positive.

Table 4. Confusion matrices: assessment result incorporating rows of trees.

		Assessment result	
		Acceptance	*Rejection*
Reference	*Correct*	66.1%	31.9%
	Incorrect	1.3%	0.7%

Consideration of all objects

		Assessment result	
		Acceptance	*Rejection*
Reference	*Correct*	66.5%	33.5%
	Incorrect	0.0%	0.0%

Consideration of only the objects being significantly concluded by trees

If only occluded road objects are considered, the increase in efficiency resulting from the use of rows of trees becomes obvious: the number of true positives increases from about 53% to more than 66% if automatically extracted rows of trees are incorporated.

Discussion of influence of rows of trees and of graph analysis. The effect of the incorporation of rows of trees into the assessment is demonstrated with the image presented in figure 10. The ten objects marked with '+' are additionally accepted road objects, compared to the results where only extracted roads have been incorporated. They all fulfill important network tasks as they link the road objects which have been accepted in Phase 1. Since the extracted road objects in conjunction with the rows of trees give enough evidence for the correctness of these important objects, they have finally been accepted. An interesting issue is related to the road object marked with '1' in the figure. This object is entirely covered by trees. According to the strategy implemented in the presented approach, this object has been rejected, because no acceptance decision is possible without any evidence given by extracted roads.

Limits of the presented graph approach are pointed out in figure 11. This scene shows that many road objects which are significantly occluded by trees, and which are connected to the network at only one node, i.e. they constitute dead end roads. One reason for this is the canal in the left part of the image. Only the objects marked with 'X' have been accepted in Phase 1. Although the rows of trees give enough evidence for the correctness of some additional road objects, those have been rejected since they do not connect the accepted roads. This example demonstrates

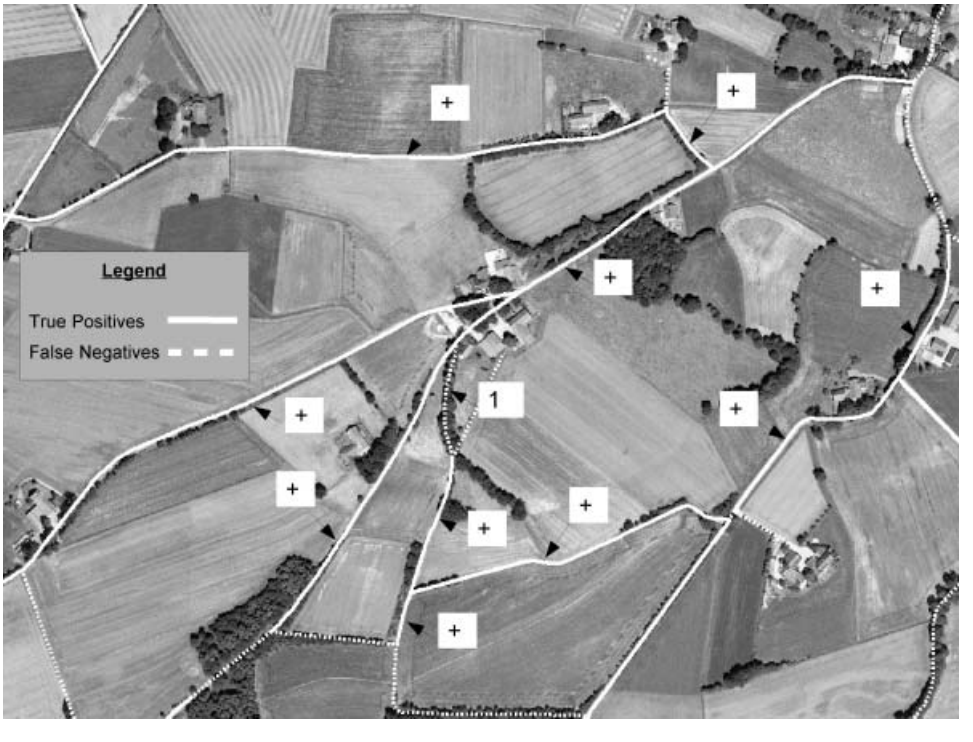

Figure 10. Example for true positives because of incorporation of rows of trees. The '+' mark objects which are additionally accepted after the rows of trees were incorporated.

Figure 11. Limits of approach. False negative decisions for dead end roads, occluded by rows of trees.

that in practical applications it is advantageous to process the scenes in a seamless manner to avoid objects at the image border.

5.3 *Discussion of false positives and false negatives*

The confusion matrices show that about 1% false positive decisions, i.e. objects being incorrect according to the decision of the human operator have been accepted by the automatic assessment algorithm. Two typical examples for those false positive decisions are presented in figures 12 and 13. The first figure shows a road object under construction, the human operator labelled it as incorrect. However, the road extraction algorithm detected a valid road object which coincides well with the given ATKIS road object. Such errors related to the actual usability of a road can only be avoided if additional information is provided to the road extraction algorithm, or to the assessment algorithm, respectively.

The second example for a false positive decision (refer to figure 13) reveals two problems inherent in the overall approach. First, the road object is correct, except for the junction area at the right image border – the geometric connection to the adjacent road is too far away from the correct medial axis. The actual offset measured in the image is about 8 m (see magnification on the right hand side in the figure). The maximum offset apart from the junction area is only 4 m and does not exceed the tolerance. Second, especially in junction areas a formal definition of the medial axis is difficult, resulting in possibly different assessment decisions. For instance, a second human operator would perhaps assume an extension of the straight line (dashed black line in the magnification), resulting in a maximum offset

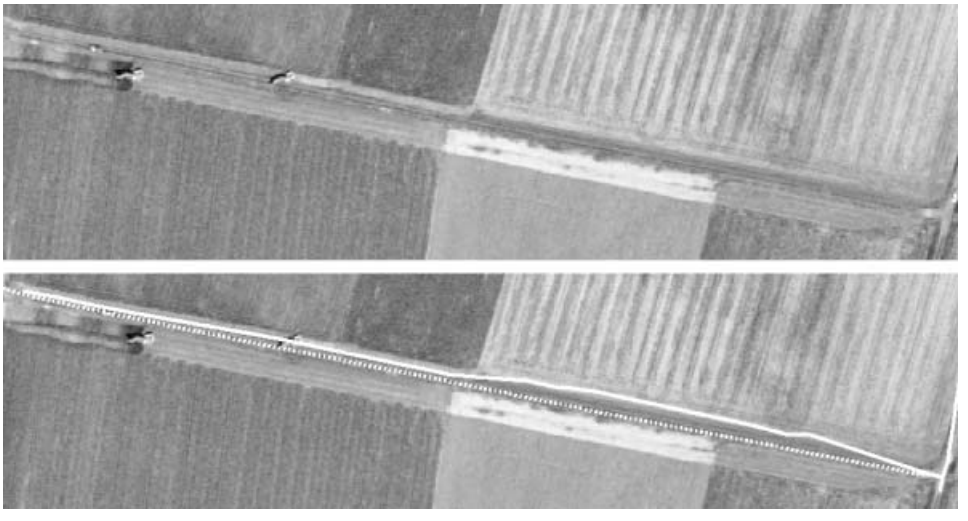

Figure 12. Example for false positive decision. Upper image: orthoimage detail showing a road under construction; lower image: superimposed extracted road (continuous line) and ATKIS road object (dashed line).

of 4 m and thus would classify the ATKIS road object as correct. Similar to the evaluation of the requirement *same shape* a certain portion of uncertainty inherent in the preparation of reference data can be observed regarding the definition of junctions.

Two main reasons lead to false negative decisions, i.e. the rejection of correct roads. If roads are occluded by local context objects, for instance by vegetation, the road extraction algorithm cannot find the object of interest. In those cases the idea presented in this paper, namely the incorporation of those local context objects into the assessment leads to a larger acceptance rate, but with the limitations described above, causing some false negatives. In other cases the road is not occluded, but the

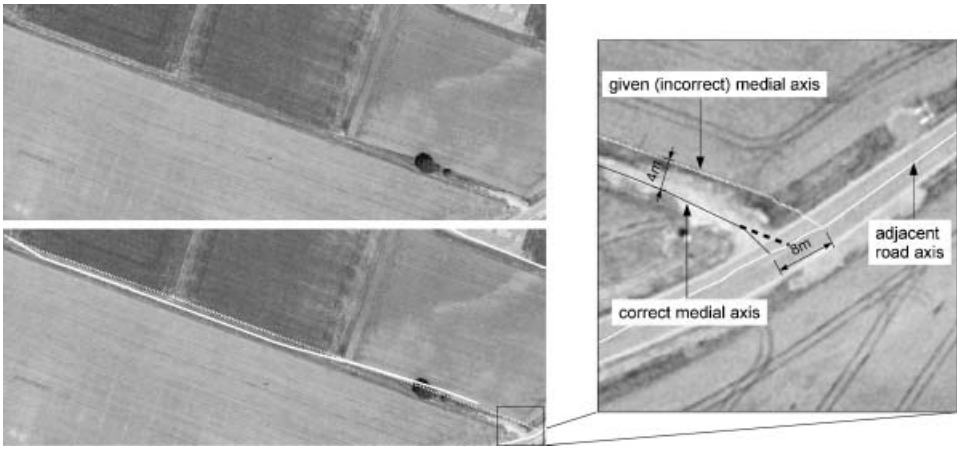

Figure 13. Example for false positive decision. Left, upper image: orthoimage detail showing a path leading to a field; left, lower image: superimposed extracted road (continuous line) and ATKIS road object (dashed line). The magnification (right) shows the correct road axis. The ATKIS road object is not correctly connected to the adjacent road at the right side.

road extraction algorithm fails to delineate the road because of weak contrast conditions. For many cases the road extraction operator used in this paper can be adapted to those cases, however, increasing the risk of additional false positives.

6. Conclusions and outlook

The presented approach to road database assessment using remotely sensed imagery uses a detailed modelling of objects and their relations, including a consistent statistical processing for the comparison of vector datasets from different origins. The evaluation of the approach demonstrates that it is worth incorporating statistic means to model the individual properties and to assess a given dataset. The explicit incorporation of context object significantly increased the performance of the method. The effectiveness of the assessment also depends on the performance of the used road extraction operator. If correct road objects are not extracted or if non-road objects which appear as roads in the imagery lead to false extractions, errors in assessment decisions cannot be avoided. However, by means of the chosen graph-based strategy which also uses context objects to explain gaps in road extraction, the number of errors is reduced considerably.

The approach to road database assessment presented in this paper is integrated in the workflow at the German Federal Agency for Geodesy and Cartography (BKG) for the automated verification and quality control of the ATKIS DLMBasis. It was shown in empirical evaluations that using our approach the amount of time which is necessary to verify a given dataset within the established workflow at BKG is reduced by a factor 3 for open landscape scenes.

The incorporation of further objects into the assessment seems to be an interesting and promising means of improvement. The relationship model can easily be extended towards new object classes. For instance, the edges of forests are not considered up to now. Similar to rows of trees they may occlude roads and therefore hamper the automatic extraction of roads. The geometric and topologic relations can be modelled similar to those for rows of trees. To integrate additional object classes is also interesting for the graph-based optimisation. In the current approach only reliably extracted roads are considered. Especially for dead end roads other object classes – for instance single man made facilities – may imply that a road object is important in the sense of connection functionality.

Additional significant improvements concern the extension of the approach regarding quality assessment in settlement areas, which requires an enhancement of the road extraction operator, the relationship model and further investigations concerning the incorporation of the existing data into automatic road extraction. Object classes to be included into the relationship model are for instance buildings, rows of buildings, and objects on the road like vehicles.

Acknowledgements

This work was funded by the German Federal Agency for Cartography and Geodesy (BKG). The support is gratefully acknowledged. We also thank the anonymous reviewers for their valuable comments.

References

AdV (Ed.), ATKIS-Gesamtdokumentation, Hannover 1989.

AdV 2007 adv-www: 2007 AdV, ATKIS-Objektartenkatalog. Available online at: http://www.atkis.de (accessed August 31st, 2007).

BLUM, H., 1967, A transformation for extracting new descriptors of shape. In *Models for the Perception of Speech and Visual Form*, W. Whaten-Dunn (Ed.) (Cambridge, MA: MIT Press), pp. 362–380.

BUSCH, A., GERKE, M., GRÜNREICH, D., HEIPKE, C., LIEDTKE, C.E. and MÜLLER, S., 2004, Automated verification of a topographic reference dataset: system design and practical results. In *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, **35**(B2), pp. 735–740.

DE GUNST, M., 1996, Knowledge-based interpretation of aerial images for updating of road maps. PhD thesis, Netherlands Geodetic Commission Publications on Geodesy, TU Delft No. 44.

DUDA, R.O. and HART, P.E., 1973, *Pattern classification and scene analysis* (New York: Wiley).

EGENHOFER, M.J. and FRANZOSA, R.D., 1991, Point-set topological spatial relations. *International Journal of Geographical Information Systems*, **5**, pp. 161–174.

GERKE, M., 2006, Automatic quality assessment of road databases using remotely sensed imagery. PhD thesis, Deutsche Geodätische Kommission Series C, No. 599.

GERKE, M., BUTENUTH, M., HEIPKE, C. and WILLRICH, F., 2004, Graph supported verification of road databases. *ISPRS Journal of Photogrammetry and Remote Sensing*, **58**, pp. 152–165.

GERKE, M. and HEIPKE, C., 2006, Automatic quality assessment of road databases in open landscape areas. In *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, **36**(4), pp. 115–120.

GIMEL'FARB, G.L., 1996, Texture modelling by multiple pairwise pixel interactions. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, **18**, pp. 1110–1114.

HU, M.K., 1962, Visual pattern recognition by moment invariants. *IRE Transactions on Information Theory*, **8**, pp. 179–187.

KRAUS, K., 1993, *Photogrammetry*, 4, Vol. 1 (Bonn: Dümmler).

LAMBERT, G. and GAO, H., 1995, Line moments and invariants for real time processing of vectorised contour data. In *Proceedings of the Image Analysis and Processing, ICIAP*, **8**, pp. 347–352 (Berlin: Springer).

LEE, D.T., 1982, Medial axis transformation of a planar shape. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, **4**, pp. 363–369.

MAYER, H., BALTSAVIAS, E. and BACHER, U., 2006, Automated extraction, refinement, and update of road databases from imagery and other data. In *Report Commission 2 on Image Analysis and Information Extraction, European Spatial Data Research – EuroSDR*, Official Publication 50, pp. 217–280.

PLIETKER, B., 1997, Automatisierte Methoden zur ATKIS-Fortführung auf der Basis von digitalen Orthophotos. In *Photogrammetric Week*, Fritsch and Hobbie (Eds) (Heidelberg: Herbert Wichmann), pp. 135–146.

ROSENFELD, A., 1986, Axial representations of shape. *Computer Vision, Graphics and Image Processing*, **33**, pp. 156–173.

RUSKONÉ, R., 1996, Road network automatic extraction by local context interpretation: application to the production of cartographic data. PhD thesis, Université de Marne-La-Vallée, Noisy-le-Grand, France.

SARDANA, H.K., DAEMI, M.F. and IBRAHIM, M.K., 1994, Global description of edge patterns using moments. *Pattern Recognition*, **27**, pp. 109–118.

SOILLE, P., 1999, *Morphological image analysis* (Berlin: Springer).

STEGER, C., 1998, An unbiased detector of curvilinear structures. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, **20**, pp. 311–326.

STRAUB, B.M., 2003, Automatische Extraktion von Bäumen aus Fernerkundungsdaten. PhD thesis, Deutsche Geodätische Kommission Series C, No. 572.

STRAUB, B.M., WIEDEMANN, C. and HEIPKE, C., 2000, Towards the automatic interpretation of images for GIS-update. In *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, **33**(B2), pp. 525–532.

WIEDEMANN, C., 2002, Extraktion von Strassennetzen aus optischen Satellitenbildern. PhD thesis, Deutsche Geod\ätische Kommission Series C, No. 551.

WINTER, S., 1996, Unsichere topologische Beziehungen zwischen ungenauen Flächen. PhD thesis, Deutsche Geodätische Kommission Series C, No. 465.

WINTER, S., 1998, Uncertain topological relations between imprecise regions. Technical Report, Fachbereich Geoinformation, TU Wien, Austria.

Appendix A: line moments

In the following the symbols according to Lambert and Gao (1995) are used. Given a two-dimensional object B, whose domain $D(B)$ is known, the moment (p, q) of order$=(p+q)$ is calculated by

$$m_{pq} = \int\int_{(x,\,y)\,\in\,D(B)} x^p y^q f(x, y)\mathrm{d}A \quad p, q = [0, 1, \ldots, \infty], \tag{A1}$$

where $f(x, y)$ is a continuous function. In this paper a linestring representation in R^2 is assumed. If an elongated object is given, the linestring representation is derived from the medial axis, reflecting the object's shape. The uncertainty inherent in the geometric position of the vertices is modelled by a Gaussian distribution and is assumed to be identical for all vertices as well as for the x and y coordinate directions, i.e. variance $\sigma_x^2 = \sigma_y^2 = \sigma_{x,\,y}^2$ for all vertices.

The line is represented by an n-side polygon with vertices (x_i, y_i) where $i=1, 2, \ldots, n$. Lambert and Gao (1995) derive the moments m_{pq} for line objects. They are given with equation (A2).

$$m_{pq} = \sum_{i=1}^{n-1} \left[\begin{cases} \sqrt{1+a_i^2} \sum_{k=0}^{q} \binom{q}{k} a_i^k (y_i - a_i x_i)^{q-k} \frac{x_{i+1}^{p+k+1} - x_i^{p+k+1}}{p+k+1} & \text{if } x_i \neq x_{i+1} \\ x_i^p \frac{y_{i+1}^{q+1} - y_i^{q+1}}{q+1} & \text{if } x_i = x_{i+1} \end{cases} \right], \tag{A2}$$

where a_i is the slope of the ith segment: $a_i=(y_{i+1}-y_i)/(x_{i+1}-x_i)$.

In case of line moments the zeroth moment m_{00} is the length of the line. The *central moments* μ, i.e. translation invariants, are calculated considering the centre of gravity $\bar{x}, \bar{y}$:

$$\mu_{pq} = \sum_{i=1}^{n-1} \left[\begin{cases} \sqrt{1+a_i^2} \sum_{k=0}^{q} \binom{q}{k} a_i^k ((y_i - \bar{y}) - a_i (x_i - \bar{x}))^{q-k} \frac{(x_{i+1} - \bar{x})^{p+k+1} - (x_i - \bar{x})^{p+k+1}}{p+k+1} & \text{if } x_i \neq x_{i+1} \\ (x_i - \bar{x})^p \frac{(y_{i+1} - \bar{y})^{q+1} - (y_i - \bar{y})^{q+1}}{q+1} & \text{if } x_i = x_{i+1} \end{cases} \right] \tag{A3}$$

The centre of gravity is obtained from the zeroth and first order moments:

$$\bar{x} = \frac{m_{10}}{m_{00}}, \quad \bar{y} = \frac{m_{01}}{m_{00}}. \tag{A4}$$

Rotation invariance is obtained applying the *method of principal axis* (Hu 1962). Here, first the orientation θ of the principal axis is calculated using second-order moments:

$$\theta = \frac{1}{2} \arctan \frac{2\mu_{11}}{\mu_{20} - \mu_{02}}. \tag{A5}$$

Two problems concerning the orientation calculation need to be considered. If the dimension of the object becomes similar in both x and y directions (circle), θ is not defined. This problem does not occur in our task, because only elongated objects are treated. The second issue is the ambiguity of θ concerning the correct choice of the major axis. The correct orientation is $\theta=\theta+n\pi/2$. Sardana *et al.* (1994) propose to use the following rules to find the correct moments:

(i) $\mu_{20}>\mu_{02}$, otherwise add $\pi/2$ to θ;
(ii) $\mu_{30}>0$, otherwise change the signs of all moments μ_{pq} with odd order $p+q$.

The object is then rotated by $-\theta$ using the centre of gravity as the point of rotation. The orthogonal transformation of rotation is given by:

$$\begin{bmatrix} x' \\ y' \end{bmatrix} = \begin{bmatrix} \cos\theta & \sin\theta \\ -\sin\theta & \cos\theta \end{bmatrix} \begin{bmatrix} x-\overline{x} \\ y-\overline{y} \end{bmatrix}. \tag{A6}$$

The central moments μ'_{pq} calculated from the rotated object are *translation and rotation invariant*.