SEMI-AUTOMATIC UPDATE AND QUALITY CONTROL OF ROAD DATABASES

J. Beyen[a], M. Ziems[b], S. Mueller[b], S. Roovers[a], C. Heipke[b]

[a]NGI – National Geographical Institute,
Abdij Ter Kameren 13, B-1000 Brussels, Belgium, {jbe, str}@ngi.be

[b]IPI – Institute of Photogrammetry and GeoInformation, Leibniz Universität Hannover,
Nienburger Str. 1, 30167 Hannover, Germany, {ziems, mueller, heipke}@ipi.uni-hannover.de

KEY WORDS: Road database, Quality, Updating, Aerial, DEM/DTM, Classification, Model, Change detection

ABSTRACT:

This paper describes a semi-automatic system for road update and quality control based on 50 cm pixel orthophotos and 3D surface models. The automatic part of the method tries to find agreements and inconsistencies between database and imagery. This is realized by an object-wise verification of all existing database records, followed by a scene-wide detection of redevelopment areas. In order to enhance the reliability, the potential success of the verification approach is measured with respect to a certain situation in the image. If a situation is not well explained by the underlying model of the approach, this image region is considered as uncertain. In the manual part of the workflow all situations that are classified as inconsistent or uncertain are used to guide a human operator to places where an update is required or where the situation is too complex for the automatic part. Experimental results achieved with data from the national Belgian road database in a test site of about 134 km² show a significant time gain, but the completeness of the updating is also lower compared to a manual process. Nevertheless, it can be recommended to alternate manual and semi-automatic updating in successive updating cycles. In order to test the second goal of our work, namely quality control of outsourced database updates, we initially limited ourselves to the use of detected inconsistencies. This was found to be too optimistic and the visual control had to be extended to all update hints to make the method useful when applied in combination with other methods for quality assessment.

1. INTRODUCTION

1.1 Motivation

As in many other areas today, also national mapping agencies like the Belgian National Geographical Institute (NGI) have to perform more and more work with less and less resources. A main part of that work concerns updating databases by manual comparison with aerial images.

In 1990 the first collection of midscale digital topographical data was accomplished with 22 stereo-operators; nowadays there remain only 8 operators to perform the updates. Moreover, whereas the first digital data collection took 18 years, the expectation is now to update the country in 3 years for transportation networks and buildings and in 6 years for the other objects. As the objects in the 3 year update cycle form 80% of the total amount of data, this means that 36% of the operators have to cover the country 5.4 times quicker than before: each operator has to cover in average 15 times more area per year than before.

In order to solve the task with the limited human resources the updating of buildings and part of the updating of the road network were outsourced to private companies. For the current update cycle 33% of the Belgian roads were outsourced. In the future about 50% will either be outsourced or be gathered through co-operation with external sources. In this scenario NGI is faced with two different problems. The first is the in-house updating; the second is the quality control of the outsourced parts. At least a partial automation of both tasks is required to reach the 3 years' cycle with the available staff.

A starting point for such an automation is given by the fact that the number of required changes per update cycle will decrease drastically as the average age of the vector data has decreased from 12 years (after the first data collection) to 3 years (after the first update). It can be estimated that only about 12% of the buildings and 3% of the roads have to be updated at every cycle. Thus, the effort of visually scanning the images for changes becomes too big compared to editing the data themselves. In the present paper we analyse the possibilities of automating the visual checking to support both tasks: the *in-house updating* and the *quality control* of the outsourced parts.

1.2 Related Work

Most recently national mapping agencies all over the world have automated the database update in two different regards. One successful approach consists in the fully automatic testing of the database consistency with respect to the data model, e.g. the NGI's validation service went very far in this direction in order to reach a perfect geometric, semantic and topological coherence of the database (Beyen et al., 2008). The second aspect is connected to the new interactive stereo-editing tools, provided with standard GIS software, that allow to trigger groups of operations or to propagate changes from an object to the adjacent ones.

Despite of such progress, the aspect of image interpretation and comparison with database is still task to human operators. However, there is a lot of active research in that field, which had started two decades ago. Already in (Mena, 2003) more than 250 different approaches dealing with road detection were listed, while a more recent overview is given in (Poullis and You, 2009). In the following we describe a few state-of-the-art approaches, representing different contributions in that research field and point out their strengths and limitations. Then we discuss their potential with respect to automatic road update before we describe our own strategy.

Road Detection: A significant group of road detection approaches model roads as linear objects in aerial or satellite imagery with a resolution of about 1-2 m, e.g. (Wiedemann and Ebner, 2000). A straight forward improvement of that model is incorporating parallel edge pairs, see e.g. (Baumgartner et al., 1999). The EuroSDR test on automatic road extraction algorithms (Mayer et al., 2006) has shown that such approaches produce already stable results for areas with homogeneous background. It turned out that nearly 85% of the roads could be successfully detected if only rural area is considered. However, different models and strategies are required for the urban context. A quite successful approach, considering urban context, was developed by Youn et al. (2008). The approach is based on orthophotos with 0.12m ground sampling distance (GSD) and LIDAR data. The underlying model focuses on the discrimination of rows of buildings, ground vegetation and roads. The presented experiments achieve a detection rate of about 80% for a challenging urban area. A disadvantage of that approach is the restriction to urban areas and to a grid-like road network structure. Other approaches like (Hinz and Baumgartner, 2003) and (Grote et al., 2012) avoid such general restrictions. Instead, they formulate a context model that considers the interaction between roads and other objects always as optional. The considered context objects are cars, buildings, trees, ground vegetation and road markings. Their experiments are based on similar input as (Youn et al., 2008), while the achieved detection ratios are a bit lower: around 75%. All these approaches rely on manually formulated road models. Recently, the use of statistical approaches, formulating the model on the basis of training data, has gained interest. Most of all, Mnih and Hinton (2010) demonstrated the functionality of such a strategy by introducing a huge amount of training data. As the training data have to cover a high variety of road appearances, the key idea is to use large parts of the existing road database for training. The experiments were carried out in $80km^2$ sub-urban test site by introducing a training area of $500km^2$. The achieved detection ratio is stated with 85%.

Road Update: Even if the given review of road detection approaches is far from complete, some lessons for updating can be learnt. First of all, the achieved detection ratios with a maximum of 85% are far from allowing a straight forward takeover into the database. This is aggravated by the fact that all the methods were originally developed for special environments and context regions, sometimes only tested for small scenes. The second big problem is the low quality of extracted road geometries, compared to manual solutions. This is especially true for crossroads and roads with varying widths. Due to these problems, only few methods for updating databases were developed so far. These methods usually circumvent the stated problems by involving a human operator. Hence, in (Klang, 1998) a semi-automatic system for an enhancement of the Swedish road database based on a comparison with satellite imagery of SPOT and Landsat is described. The approach detects the position of road junctions within a tolerance radius around the position, indicated by the database. Based on this result the nodes are used as seed points for an active contour model which is applied to every road object of the database. Finally, a comparison of the extraction result and the corresponding database object provides the human operator with a number of potential objects for the updating process. The system was extended in relation to the task of the National Topographic Database of Geomatics Canada (Fortier et al. 2001). Gerke et al. (2004) describe a semi-automatic approach for quality control in rural area. They apply the approach of Wiedemann and Ebner (2000) for image regions, indicated by the database and then compare its outcome with the database. All the stated agreements are skipped for the manual quality control, which leads to an automation ratio of about 60% in rural areas. In (Gerke and Heipke, 2008) the approach is extended by extracting rows of trees to explain possible occlusions of roads. Poulain et al. (2010) describe a method that applies high resolution SAR and optical images for an automatic update of a road database in urban context. In this work, for each road object, features are extracted within a region around the position, indicated by the database. The features reflect different properties of roads but also those of typical urban context objects. Then, the authors consider the two classes correct and incorrect in a classification step. In an experiment the test delivered correct decisions for more than 90% of the objects. As this strategy is not able to detect omission errors the authors add a second step where road candidates are extracted from all over the image and verified by their proximity to the road network that was verified in the first step. This second step ends up with a detection rate of about 60%. The work, described in the present paper is developed in co-operation with the WiPKA project that was initiated in 2000 by the federal mapping agency of Germany for the task of quality control of topographical data (Busch et al. 2004; Helmholz et al. 2012). Within the project a method for road database verification on the basis of aerial imagery was presented in (Ziems et al., 2011). Similar to (Gerke and Heipke, 2008) and (Poulain et al., 2010) a hypothesis test is carried out for all entities of the existing database in order to detect inconsistencies with the up-to-date imagery.

1.3 Basic Concept

The goal of the present paper is to describe a system that allows firstly an effective and reliable *in-house updating* and secondly a nearly fully automatic *quality control* of the outsourced parts. In accordance with the review of the related work we use a semi-automatic approach in which we introduce the basic techniques from the WiPKA-project to detect possible inconsistencies. These inconsistencies are used as hints to guide the human stereo-operator to areas where updating is needed. Thus, the time for visually scanning the images can be reduced. Besides the database to be updated and the images height information in form of a normalised digital surface model (nDSM) is used as input data. Any changes to the database are still the task of the human operator.

In order to cover both of our tasks, we select two different configurations of the WiPKA road verification approach. The first configuration is directed to generate hints for in-house updating. As a high reliability is the main request for that task, all road objects not explicitly verified are checked by the operator. The second configuration is directed to the quality control task. As here a high automation is the main request, the operator only checks clearly detected inconsistencies. The verification approach only considers existing database entities, which requires a second method to detect also new roads. Due to the high difficulty of detecting roads in arbitrary context we detect simply new development areas instead.

The focal point of the paper is the analysis of the proposed approach in terms of practical applications. Thus, we carried out a large test under production conditions in which the results of a human stereo-operator, supported by the proposed method, are compared to those of an operator without support. In doing so to answers to the following three questions are sought:

- How many changes are missed?
- How big are the missed changes?
- How much time can be gained?

2. DATA PRE-PROCESSING

Three kinds of input files have to be prepared for the automatic component: vector data, nDSM and orthophotos. All input data have to be georeferenced in the same co-ordinate system, in our case the Belgian Lambert2008.

Vector data: Basically, 3 different feature types in our database provide relevant information about the road network. These are road segments, dirt road segments and brunnels (bridges or tunnels). They are combined into a single road layer, which is then exported in Shapefile format. The combined layer represents a single road class with the attributes type, width and brunnel, where a single road object is characterized by a single attribute entity. In accordance with the standard GIS definition, nodes with a degree greater than 2 are always endpoints of an object. Furthermore, a second Shapefile layer is exported that gives a complete scene description with respect to the two classes *developed* and *undeveloped*. While the first class is computed by buffering roads with 15m and buildings with 50m width, the second class is simply the complement of the first.

nDSM: A digital surface model (DSM) and a digital terrain model (DTM) are generated with the commercial Match-T DSM software from stereo images with 0.3m GSD. Since normally only DTMs with a grid size of 20m DTM are generated, the DTM process needed to be adapted to compute a 1.5m DTM/DSM from which the nDSM was calculated. For this computation only points were used for which the correlation was "sure". Finally, the resulting nDSM is transferred from ASCII into a floating point image in geotiff-format.

Orthophotos: At NGI colour-orthophotos (RGB or IRRG) are stored in geotiff-format in a database in tiles of 2 km x 2 km with 0.5m GSD. The orthophotos are computed from aerial images with 0.3m GSD. For the automatic analysis no further pre-processing of the orthophotos was required, except for mosaicking the orthophotos to cover the whole target area, which allows computing bigger areas in one cycle.

3. AUTOMATIC COMPONENT

The automatic component is designed as a black box, running on a LINUX system. The interface only requires the pre-processed data. Internally, a twofold method is applied. The first method verifies all road objects by analysing the corresponding image regions. Thus, changed or removed roads can be detected. In contrast, the second method indentifies new roads by detecting new development areas. In the following both methods will be briefly described, while a comprehensive overview is given in (Ziems et al., 2011).

3.1 Object-wise road verification

For each road object a corresponding sub-scene is extracted from the image/nDSM mosaic that contains the road and its surrounding area up to a depth $d = 15 \cdot road\ width$ (see Figure 1). In each sub-scene the database hypothesis is tested in terms of geometry and the attributes type, width and brunnel. As the visual properties of roads can be very different, a test with general applicability has to consider many possibilities. To circumvent a complex analysis a bundle of relatively simple hypothesis tests with different strengths is applied instead.

Figure 1: Considered sub-scene with road hypothesis from the outdated database.

The applied hypothesis tests rely on road detection/verification approaches from recent literature. The approaches are selected so that different model assumptions are reflected. The underlying models describe roads e.g. as lines in homogenous context, as areas with discriminative spectral and textural properties to buildings and parks and as 3D valleys in the nDSM.

The correct combination of all these approaches leads to a more general solution. However, due to the restrictive model assumptions, each algorithm only provides reliable decisions for specific situations. Finally, the applicability of an algorithm is connected to the compliance of the model assumptions to the actual appearance of a sub-scene. Therefore, each sub-scene is analysed in order to determine a confidence value C with $0 \leq C \leq 1$ that reflects the degree to which the situation corresponds to the optimal situation for an approach. Then, the decisions of all algorithms are combined in a decision level fusion process in which the confidence values control the impact of a single decision on the final result. The combination is based on the theory of Dempster-Shafer (Dempster, 1968 and Shafer, 1976), which provides tools for dealing with uncertain and incomplete knowledge about the statistical properties of the data.

Figure 2: Object-wise verification method output
(correct=green, incorrect=red, uncertain=yellow).
Left: detected database error.
Right: hardly visible but correct road.
Bottom: exemplary output.

After applying Dempster's rule any decision is connected to an uncertainty mass. On this basis three cases for a database object are considered. Firstly, the automatic analysis highly supports the correctness; secondly it highly supports the incorrectness and thirdly, the automatic analysis is inappropriate to provide any reliable decision. Thus, the verification ends up with a traffic light decision for each road object, where undecidable cases are depicted in yellow, for an example see Figure 2.

For the experiments presented in section 5 nine different road verification algorithms are used. As a side note, we mention that the separate computation of each algorithm and road object allows a broad parallelization, which keeps the computational times feasible also for large databases.

3.2 Development area detection

This second method aims at detecting new roads and omission errors in the database. Compared to the first method, the object-wise road verification, this task is much more challenging, as less prior information is available. We circumvent the problem of road extraction by simply detecting new developed areas, where we assume the majority of the new roads to be situated. Thus, this second method verifies all the regions that are denoted as *undeveloped* in the database. To achieve this aim, the support vector machine classifier (Vapnik, 1998) is applied to separate the two classes *developed* and *undeveloped*. The feature vector is defined for each pixel by the spectral mean and variance of the gradient image within a 15m neighbourhood. In addition to the available image channels the nDSM is used for the feature computation. The area-wide classification result is compared to the pre-processed landcover layer. Then, all contradicting regions that are classified as *developed* are computed. In order to filter out irrelevant objects, we apply a minimum threshold of 1500 m^2 for the region size. The resulting regions are vectorized and introduced in a GIS environment as areal hints for new roads (see Figure 3).

3.3 Outcome

Both methods deliver object lists in Shapefile format. The list, computed by the *object-wise verification tool*, delivers hints for incorrect and hardly detectable roads (red and yellow). The second list, computed by the *development detection tool*, delivers regional hints for new development areas. The hints of the second list are always located in previously undeveloped regions, while new roads in already developed areas are often associated with a rearrangement of roads and thus part of the first list.

Figure 3: Development area detection
Left: Pre-processed landcover layer
Right: classification result
(developed= red, undeveloped = green)
Bottom: exemplary output
(road data = yellow, update hints = red)

4. INTERACTIVE POST EDITING

The output of the automatic component is transformed into the GeomediaAccessWarehouse format from Intergraph and introduced into ISSG, which is the standard stereo-editing tool at NGI (see Figure 4). For the interactive post editing the imagery and relevant road layers are overlaid with the computed hints as thick red lines. In order to guide the operator quickly to the areas where update is required, the operator picks the hints, one by one, from an overview mono view (see Figure 4 right). A picked hint is then centred in the close-up stereo view for detailed investigations. The hint is deleted and the possibly required editing is carried out on original road layers in the same way as for a completely manual update. The operator then proceeds the editing around the just deleted hint as long as there are changes connected to the previously edited objects. When no further changes can be spotted the operator picks the next hint.

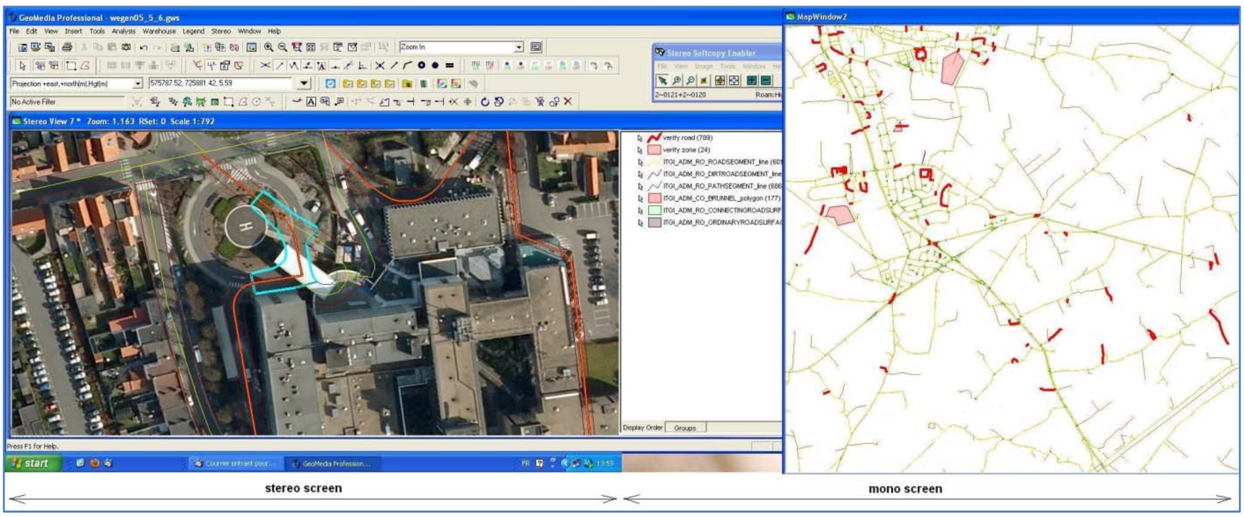

Figure 4: Visualisation of hints and other objects in mono and stereo during the interactive post check

5. EXPERIMENTS

As introduced in section 1.3 in this paper the practical usability of the method for mapping agencies is explicitly analysed. In contrast, the scientific contribution of the proposed approach is investigated in (Ziems et al., 2012) on the basis of a benchmark test.

As the performance of automatic road detection depends heavily on the appearance of roads, we selected a 134 km² area situated at the coast to evaluate our strategy. The region contains a whole town but also suburban and rural parts. In summary 6496 road objects with 686 km total length are situated in that area. The precondition for the evaluation is as follows: we have vector data for roads and buildings with the timestamps 1991 and 2008. The latter was updated at NGI on the basis of the former one by a manual comparison with RGB imagery from 2008. The update process was widely documented so that the time needed for the process is exactly known. Furthermore, IRRG imagery from 2011 is available. The manual update of the 2008 data with the 2011 images is not yet available for evaluating the method. In the following we present the experiments that are oriented to the both tasks *in-house updating* and *quality control*.

5.1 In-house updating

In accordance with the manual update *1991-2008* we applied the proposed method with the 1991 vector data and the 2008 imagery. As introduced in section 1.3 we use *all* unverified road objects (yellow + red) plus all detected new development regions as hints to guide the human operator. The automatic component resulted in 813 hints, of which 110 were denoted as inconsistencies (red), 679 were denoted as unclear (yellow) and 24 were denoted as new development regions. Then, a human operator was asked to update the road network on the basis of the computed hints. In order to exclude differences in the results due to different human factors, the operator for the test was the same as the one who had carried out the manual update.

	manual	semi-automatic	diff.
required time (days)	42.00	8.25	-80%
changed road segments	3987	1453	-64%
missed new road segments	47 (3km)	200 (16km)	
missed new road segments (%)	7.3	31.0	+23.7%

Table 1: Evaluation results: manual vs. semi-automatic update

Table 1 displays the results of the test compared to the manual solution. The gain in time of 80% is clear and made it worthwhile to check the quality of the results. At first we found that the number of objects that have really been updated is very different: 64% less edited road segments in the semi-automatic solution. However, these edits are mainly related to geometric corrections of the 1991 data and not to changed roads. In order to filter these corrections from the real changes, we computed the *missed new road segments* from the spatial differences of the two solutions considering a tolerance of ± 4m. Then, we visually checked the computed differences and sorted out all objects that did not really change. An example is given in the left part of Figure 5, where a geometrically corrected junction violates the ± 4m threshold. The remaining objects are compared to the ground truth, which includes 645 new objects with a total length of 34.6km. Thus, a large proportion of edits is related to geometric corrections that will not be necessary for the following update cycles. Hence, the determined gain in time is significantly related to the omitted geometric corrections in the unseen parts.

Figure 5: Missed changes of the proposed update method. Left: small geometric correction. Right: important change

The semi-automatic update missed 200 roads, which all have the status of local or private road. Most omissions are related to small building projects in the vicinity of already existing built-up areas, cf. example in Figure 5. The longest 3 missing road segments are 847m, 624m and 283 m; 38 other missing road segments are more than 100m long.

We also tested the automatic component with the *2008-2011 update scenario* for which we have no manual reference dataset until now. Thus, we can only assess the outputs of the automatic component. The absolute number of computed hints is smaller by 436, compared to the former result of 813. Note that the true hints diminished more than the false hints, so that the 4 times better up-to-dateness of the data does not lead to 4 times less hints. A part of the reduction can also be attributed to the availability of the infrared channel and the better geometric quality of the vector data used as input. It is clear that starting from the second update cycle, thanks to the lower number of hints, the semi-automatic update can be expected to offer an even bigger gain in time compared to a classical update.

5.2 Quality control

For the second task, the manually updated database has to be compared with the imagery. As can be seen in Table 1, the human operator forgot to add 47 objects during the manual update. These errors can be detected if both kinds of hints (yellow + red) are checked. However, in (Ziems et al., 2011) it was shown that the number of false alarms is significantly lowered if only the clearly detected inconsistencies (red) are considered. Thus, we had hoped to be able to restrict ourselves to the verification of the 'red' objects and to avoid spending time on checking false alarms, but this was too optimistic. In fact, only 23 of the 47 missing road objects were detected by the 'only red' configuration, which is too little to make a general statement about the data quality.

6. CONCLUSIONS

The method presented in the paper verifies road database entities automatically by comparison with up-to-date imagery and furthermore investigates all previously undeveloped regions for new roads. In combination with a human operator the object verification turned out to be very reliable. In contrast, the second part failed in detecting smaller build-up projects and thus led to an increased number of omission errors.

The presented experiments investigated two potential tasks. First, the in-house updating, where a human operator works along the list of hints produced by the automatic component: The time gain of that guided update was determined as 80% compared to a completely manual process, while about 24% more new roads were not detected. The latter is not acceptable, even if the errors do not concern main roads. Thus, the guided update as used in the experiments is not yet good enough for systematic use but it may be interesting to be performed every second update cycle. Then, the time gain would still be significant and the lowered data quality can be recovered 3 years later. In a certain sense alternating semi-automatic and manual updating can even be recommended because the different ways of looking at the data lead to complementary corrections.

For the second task, the quality control of outsourced data, we intended to only use clearly detected inconsistencies to quickly check the data with minimal human efforts. It turned out that if the data quality is good only very few errors are found. It is better to spend a bit more time for checking all available hints, but even then the result was found to be incomplete. Hence, we recommend the use for the second task in combination with visual investigations of randomly selected map samples and continued database consistency assessment. Of course, the use is only recommended if the update providers do not apply a comparable semi-automatic method for their update.

The presented results show that in the present stage of development an automatic support for updating road databases is basically realistic. In particular we identified potential improvements for the proposed strategy to achieve the goal of practical usability. First of all the omission detection has to be improved e.g. via a more detailed class description of the development areas, i.e. replacing the class ‘developed area’ by the two classes ‘buildings’ and ‘sealed ground surface’. In second place further NGI specific attributes and database rules could be considered to better deal with exceptions and generalization effects, in order to diminish the false alarms. Finally, the simple interactive post editing strategy could be replaced by a more specialized human-machine interface, e.g. by the sequential editing tool described in (Beyen et al., 2008).

REFERENCES

Baumgartner, A., Steger, C., Mayer, H., Eckstein, W., Ebner, H., 1999. Automatic road extraction based on multi-scale, grouping, and context. *PE&RS* 65(7):777-785.

Beyen, J., Henrion, J., Van de Velde, S., 2008. Selection and customization of an integrated Digital Photogrammetric Workstation + GIS configuration and optimization of the interoperability within the workflow for updating the Belgian topographical reference database. In: *IntArchPhRS vol. XXXVII, part B2, Beijing, China*: 757-762.

Busch, A., Gerke, M., Grünreich, D., Heipke, C., Liedtke, C. E., Müller, S., 2004: Automated Verification of a Topographic Reference Dataset: System Design and Practical Results. *IntArchPhotogramRemSens SpatialInformSci* 35/B2: 735-740.

Dempster A. P., 1968. A generalization of Bayesian inference. - *Journal of the Royal Statistical Society* (30) 205-247.

Fortier, M. A., Ziou, D., Armenakis, C. & Wang, S., 2001: Automated correction and updating of road databases from high-resolution imagery. – *Can. J. of Rem. Sens.* 27 (1): 76–89.

Gerke M., Butenuth M., Heipke C.,Willrich F.,2004. Graph supported verification of road databases. *ISPRS Journal of Photogrammetry and Remote Sensing*, 58(3-4), pp. 152-165.

Gerke, M., Heipke, C., 2008. Image based quality assessment of road databases. *IJGIS* 22 (8): 871-894.

Grote A., Heipke C., Rottensteiner F., 2012: Road Network extraction in suburban areas, Photogrammetric Record. The Photogrammetric Record 27(137):8-28.

Helmholz, P., Becker, C., Breitkopf, U., Büschenfeld, T., Busch, A., Grünreich, D., Müller, S., Ostermann, J., Pahl, M., Vogt, K., Ziems M., Heipke C. 2012: Semi-automatic Quality Control of Topographic Datasets. *Accepted for publication In PE&RS.*

Hinz, S., Baumgartner, A., 2003. Automatic extraction of urban roads from multi-view aerial imagery. *IJPRS* 58(1-2): 83-98.

Klang, D., 1998: Automatic detection of changes in road databases using satellite imagery. *IntArchPhotogramRemSens SpatialInformSci* 32/4: 293-298.

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

Mena, J., B., 2003: State of the art on automatic road extraction for GIS update: a novel classification. - *Journal Pattern Recognition Letters archive* 24 (16), 3037-3058.

Mnih, V. and Hinton, G., E., 2010. Learning to detect roads in high-resolution aerial images. In *Proceedings of the 11th European conference on Computer vision: Part VI* ECCV'10, Springer-Verlag, Berlin, Heidelberg, 210-223.

Poulain, V., Inglada, J., Spigai, M., Tourneret, J.,Y., Marthon, P., 2010: High resolution optical and sar image fusion for road database updating. *IGARSS 2010*: 2747-2750.

Poullis, C. and You, S., 2010: Delineation and geometric modelling of road networks. *IJPRS* 65(2), 165-181.

Shafer, G., 1976: The Mathematical Theory of Evidence, *Princeton University Press.*

Vapnik, V. N., 1998: Statistical Learning Theory. Wiley, N.Y.

Wiedemann, C., Ebner, H., 2000. Automatic completion and evaluation of road networks. *IntArchPhotogramRemSens SpatialInformSci* 33(B3/2): 979-986

Youn, J., Bethel, J. S., Mikhail, E. M., Lee, C., 2008: Extracting urban road networks from high-resolution true orthoimage and Lidar, *PE&RS* 74(2): 227-238.

Ziems, M., Beyen, J., Mueller, S., Roovers, S., Heipke, C., 2011: Multiple-Model based update of Belgian reference road database. In: *IntArchPhRS* vol. XXXVIII, part 4, Guilin, China, 2011, pp. 73-80.

Ziems, M., Breitkopf, U., Heipke, C., Rottensteiner, F., 2012: Multiple-model based verification of road data. *Accepted for publication in IntArchPhRS ISPRS Melbourne 2012.*