Photogrammetric Week '15 Dieter Fritsch (Ed.) Wichmann/VDE Verlag, Belin & Offenbach, 2015

From Nationwide Point Clouds to Nationwide 3D Landscape Models

George Vosselman, Sander Oude Elberink, Enschede
Marc Post, Zwolle
Jantien Stoter, Apeldoorn/Delft
Biao Xiong, Enschede

ABSTRACT

While an increasing number of cities is using 3D geo-information, applications like infrastructure planning and the assessment of environmental impact by noise, wind and air pollution require 3D geo-information over areas covering multiple municipalities. To enable such applications a national 3D landscape model has been produced in the Netherlands by automated fusion of the national 1:10,000 scale topographic database TOP10NL with the national elevation data AHN-2. This paper reports on the design and results of this project merging 15 million 2D object with a point cloud of 640 billion points. At the scale of 1:10,000 LoD1 modelling of buildings was considered sufficient. For applications at a larger scale LoD2 models will be required. A strategy for LoD2 building modelling based on roof topology graphs combined with automatic as well as interactive error correction has been developed.

1. INTRODUCTION

In the past years an increasing number of cities acquired 3D city models to support decision making. The 3D data acquisition was often triggered by the need to support a specific task, sometimes only for a local project area, and was typically not integrated in the geo-information infrastructure of a city. Such rather ad-hoc collection of 3D geo-information is hampering the reuse of the information in projects at a regional scale.

In the Netherlands the 3D pilot brought together around 70 professionals from government agencies, industry and academia to stimulate and coordinate the developments on 3D geo-information. Amongst others this resulted in a national 3D standard (Stoter et al., 2013), 3D implementation specifications (Stoter et al., 2014) and a list of use cases for 3D geo-information (Geonovum, 2011). Several use cases like the planning of infrastructure and the assessment of environmental impact by noise, wind and air pollution require 3D geo-information over areas covering multiple municipalities or even provinces. To address these needs the Kadaster took the initiative to create a national 3D landscape model. The national 3D landscape model contains LoD0 representations for all topographic objects available in the 2D topographic database (scale 1:10,000). Vertical surfaces are inserted in case of height jumps at the boundary of adjacent objects. For buildings LoD1 models were generated. The model was derived by combining the 2D database with the nationwide point cloud captured for the production of the national elevation model (Oude Elberink et al., 2013). This paper reports on the considerations and efforts required to apply the earlier developed methodology (Oude Elberink and Vosselman, 2009) to the large datasets of a whole country. We discuss various implementation issues and results. In the second part of this paper we also describe the efforts to automate the production of building models at LoD2 based on recognition of target shapes in roof topology graphs combined with automated recognition and correction of building modelling errors.

2. THE NATIONWIDE LOD1 MODEL 3D TOP10NL

2.1. Datasets

2.2.1. TOP10NL

The TOP10NL dataset is a national object oriented dataset at a 1:10.000 scale. The dataset is open data and maintained by the Dutch Kadaster. For the data fusion with the national point cloud it was the most suitable dataset as it is produced in a homogeneous way for the whole country. A uniform countrywide national dataset at approximately 1:1000 scale (produced by maintainers of public space like municipalities) will be available from 2016 onwards. Considering the potential applications the 1:10,000 scale was considered appropriate for the first national 3D topographic dataset. At this scale object boundaries have an accuracy of 1-2 m. Object shapes may also be (slightly) generalised, e.g. buildings separated by less than 2 meters are merged. The TOP10NL contains around 15 million objects. All objects of the classes Land Use, Water, and Road together provide a complete 2D space partition. Buildings are modelled independently (i.e. as additional layer) and may overlap with objects of the three mentioned classes. In the case of bridges and multilevel road crossings, object polygons are stacked where an object is provided for each height level. The height levels do not contain a real height value, but only indicate the sequence in which the crossing objects are stacked.

2.1.2 AHN-2

The AHN-2 is the second version of the national elevation model obtained by airborne laser scanning. The point cloud captured between 2009 and 2012 contains about 640 billion points on the land surface of 33,000 km^2. While most surveys aimed at a minimum of 8-10 points/m^2, strip overlaps and multiple echoes in vegetation led to an average point density of almost 20 points/m^2. The accuracy is very high. Systematic height and planimetric errors are typically below 5 and 15 cm respectively. Standard deviations are below 5 cm in both height and planimetry (Vosselman, 2012). Data providers were asked to classify the point data into ground and non-ground. As of March 2014 the AHN-2 is available as open data.

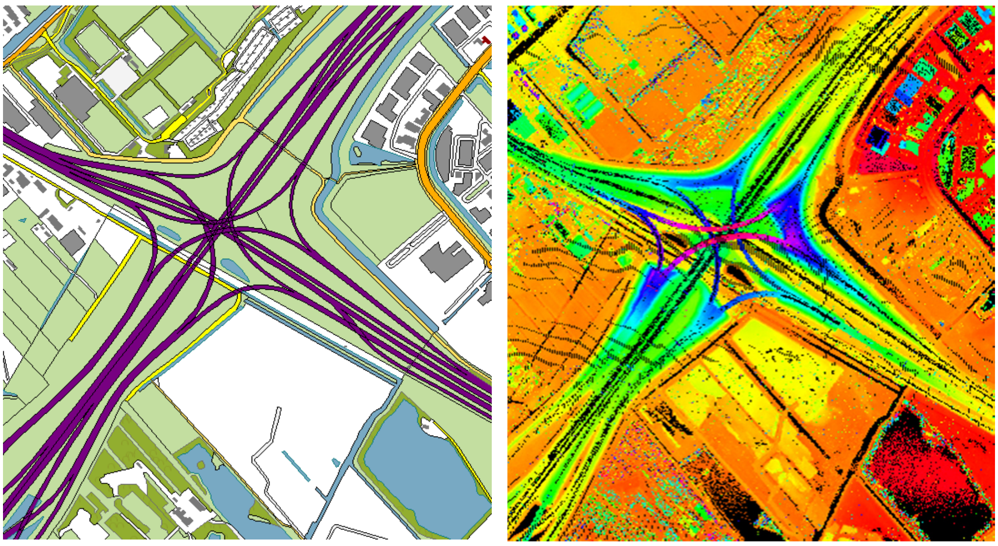

Figure 1: Example of the TOP10NL and (height coloured) AHN data.

2.2. Specification

The goal was to obtain a continuous 3D surface description of the landscape partitioned into Land Use, Water, and Road objects corresponding to their 2D object descriptions and to put on top of that surface the 3D objects representing buildings and vegetation. The 3D landscape should have no gaps with the exception of passages underneath bridges or in multi-level road crossings. Point and line elements of the TOP10NL have not been taken into account so far.
Buildings were modelled at LoD1, i.e. as prismatic solids with horizontal roof surfaces. The AHN point cloud clearly shows the roof shapes and would allow modelling at LoD2 level. However, as the objects boundaries in the TOP10NL have been captured with a standard deviation of 1-2 m in planimetry and may also have been generalised, the combination of a very accurate roof shapes with far less accurate wall locations would lead to unrealistic and visually unattractive 3D building models. At the scale of 1:10.000 LoD1 models were therefore preferred.

2.3. Data fusion

The AHN point cloud contains more than sufficient detail for 3D modelling at a 1:10.000 scale. To reduce computation times, the point cloud was first reduced to a point density of 3 points/m^2. The fusion of the point cloud with the map polygons proceeds in two phases. First, we model the 3D surface of each object independently. In the second phase the surface heights at the object boundaries are adjusted to enforce a logical transition between neighbouring objects.

2.3.1. Modelling individual objects

The strategy to model an object's surface varies with the object class (Oude Elberink et al., 2013). For building and water objects a flat horizontal surface is derived. For roads a smooth surface is generated by triangulating the heights determined for the road edges. I.e., we do not model any non-linear variation in height across the road. For terrain and vegetation (canopy) surfaces local height variations are allowed and modelled by a constraint triangulation of all lidar points inside a polygon together with the surface heights determined for the nodes at the boundary. To determine the local surface heights at object boundaries (for roads, terrain, and vegetation) we not only need to determine heights at the nodes of the 2D map polygons, but also for many locations in between those nodes as lines that are linear in 2D are not always linear in 3D. Therefore additional nodes are inserted into the 2D polygons (cf. Figure 2). At each of these nodes a local neighbourhood is defined by a circle with a radius of 5 m. The lidar points within the intersection of the circle with an object polygon are used to determine the local object height. Because of the 1-2 m noise in the map polygons not all selected points will necessarily belong to the processed object. To obtain a robust estimate we first segment the point cloud into smooth segments. The points belonging to the segment with the largest number of points within the initial selection will be used to fit a local plane. The height of this plane at the location of the node is taken as the local object height. The selection of the right points is also supported by the available classification of the points into ground and non-ground points. Hence, for modelling the road and terrain objects the selection will be constrained to points with a ground label.

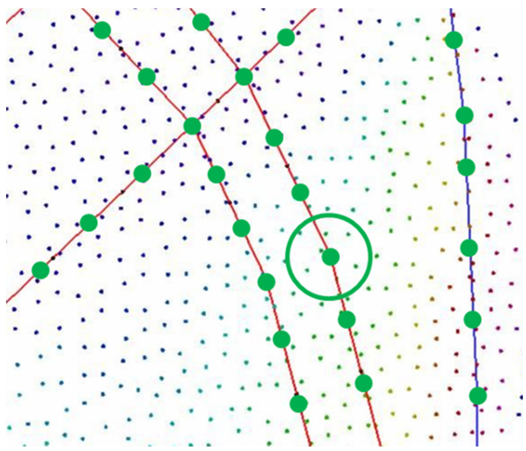

Figure 2: Local object heights are determined at locations of the 2D polygon nodes as well as many nodes inserted in between.

2.3.2. Modelling at object boundaries

With the above procedure at each node heights are determined for every object polygon through that node. This results in two or more heights for every node. The determination of the final height of a node depends on the classes of surrounding objects as well as on the difference between the derived heights (Oude Elberink et al., 2013). If the differences are small (a threshold of 20 cm has been used), the two or more adjacent objects should have the same node height. If objects have the same class label the average height is taken for road and terrain surfaces. If the class labels differ, the height of e.g. water and road surfaces will prevail over the height of adjacent terrain objects as the water and road surface heights are likely to be more accurate. If the object surfaces are clearly separated in height, e.g. between a quay (road) next to a water surface, a vertical wall polygon needs to be inserted in order to avoid gaps in the 3D surface. In the case of fly-overs large height differences between two road surfaces or a road surface and the terrain below will occur. In this case both surfaces keep their own height and a road thickness is assumed for the higher road surface to create a consistent 3D landscape model. In this way processing rules have been established for all combinations of classes. The reader is referred to (Oude Elberink et al., 2013) for further details.

2.3.3. Modelling using tiled datasets

This method for 3D landscape modelling was initially developed for datasets that can be kept entirely in the computer memory. Applying the data fusion to two national datasets with millions of objects and billions of lidar points requires a partitioning of the datasets in tiles. This partitioning should not only solve the memory problem, but also enable an independent processing of the data in tiles such that many tiles can be processed in parallel.

In a pre-processing step the AHN-2 point cloud was split into tiles of 1 km^2. For each tile all TOP10NL polygons that partially or entirely overlap with the tile are selected. When an object polygon extends over multiple tiles we need to ensure that the 3D object surfaces in the adjacent tiles seamlessly connect to each other. This is achieved by reconstructing an object surface for the part in the currently processed tile as well as a zone into the neighbouring tiles. After triangulation of the resulting surface, only the meshes with the mesh centre in the current tile are kept as the landscape model for this tile (Figure 3). In neighbouring tiles the same process is run independently and again reconstructs the object surface in a zone around the tile boundary. As we use the same lidar data and object polygons in all neighbouring tiles, we simply repeat the reconstruction of the surfaces in zones around the tile boundaries. This redundancy in the computation of the heights along tile boundaries is required to enable an independent and parallel processing of all tiles. This strategy results in seamless transitions for objects with locally determined surface heights (roads, terrain, and vegetation). Buildings and water surfaces are, however, modelled by a horizontal surface at the average height of all lidar points within a polygon. For buildings this does not constitute a problem as building polygons do not extend beyond 1 km and are therefore always completely contained in the 3x3 km area of the currently processed and adjacent tiles. The height of a certain building as computed in different adjacent tiles will therefore always be identical. This is not the case for water bodies extending over three tiles in one direction. Here the current processing strategy may result in gaps between water meshes reconstructed

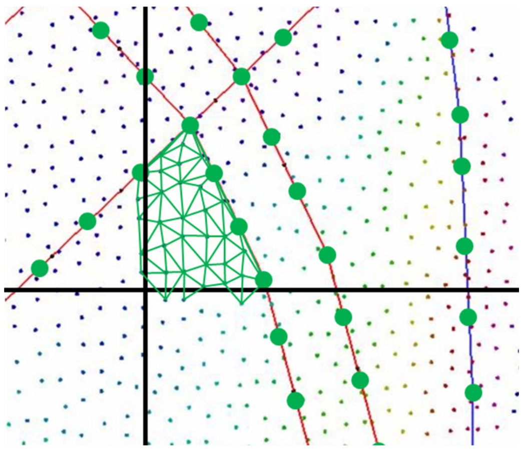

Figure 3: In each tile object surfaces are modelled by triangles up until the tile boundary.

in one tile and a neighbouring tile. Very large water bodies like the IJsselmeer or Wadden sea are not modelled at all since no laser scanning flights are conducted over these areas.

2.4. Implementation and results

The 3D landscape model for land area of the Netherlands has been produced in 30.000 tiles of 1 km^2. With an average processing time of 2.5 hours per tile it would take about 8.5 years on a single CPU to complete the national 3D model. Obviously, parallelisation is necessary. This has been done on the SARA national supercomputer in Amsterdam. Data was prepared by the Dutch Kadaster. The software of the University of Twente and the data were transferred to the SARA computer by the company Geodan. By using 100 processors on the average all tiles were processed in one month time. The analysis of the initial run showed that 90% of the tiles were processed successfully. Various reasons were identified for failure in the other tiles. It was noticed that processors sometimes ran out of memory when polygons contained many points. A 64-bit version of the software should prevent this in future runs. As the point density in the tiles varied with the companies that were contracted for the lidar surveys, the frequency of the memory problem showed a correlation with the data provider. One instance was noticed of a crash of the supercomputer as the processing of various tiles at random locations in the Netherlands all failed at the same point in time. Finally, data of some tiles was simply lost in the data transfer in between the Kadaster and SARA. After eliminating some obvious mistakes the tiles without output have been processed again, leading to the current completeness of 97%.

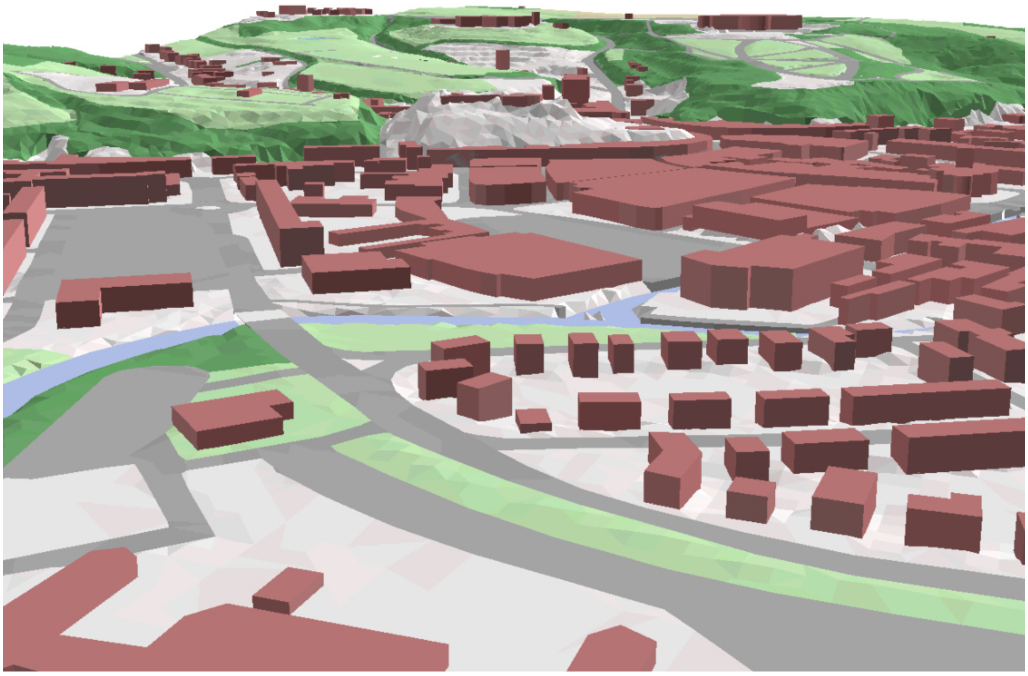

Figure 4: Screenshot of a small part of the 3D TOP10NL.

No larger quality control of the reconstructed 3D landscape has been performed so far. Based on visual inspection the quality of most tiles seems to be satisfactory (Figure 4). As mentioned above larger water surfaces were not properly modelled at tile boundaries. Points on building walls that were incorrectly classified as terrain points led to some local errors in the terrain surface. Furthermore, some extreme peaks were observed in forest surfaces. These may have been caused by errors in the range measurement or by the occasional reflection of a laser pulse on a bird. While these points were correctly classified as non-terrain, it is clear that not all non-terrain points in a forested area can automatically be considered as a reflection in the tree canopy.

Further quality analysis and improvement of the modelling will be needed. So far the project to produce the 3D TOP10NL mainly served to assess the feasibility of executing such a large computational project. Priority was also given to making the 3D model available as open data as soon as possible.

3. TOWARDS A NATIONWIDE LOD2 MODEL

For applications like urban planning, solar energy potential assessment, and virtual tourism the modelling of buildings at LoD1 no longer suffices. To model the shapes of building roofs (LoD2) is much more challenging as it needs to deal with the modelling of sometimes very small roof faces as well as the reflection of laser pulses on objects like antennas, chimneys, and dormers. Haala and Kada (2010) provide an overview on various approaches to building modelling. One class of methods that seeks a balance between data and model driven approaches was introduced by Verma et al. (2006). They recognised shapes of building parts by matching their so-called roof topology graphs against the graphs of standard roof shapes collected in a library. For LoD2 modelling of large areas it is necessary to achieve a high success rate of automatically modelled buildings. In the remainder of this paper we describe how the method by Verma et al. can be adapted to increase the variety of buildings that can be modelled with this approach. We also present a strategy on how to automatically detect and correct some of the remaining modelling errors.

3.1. Datasets

For the experiments described below a point cloud with a density of 20 points/m^2 was used. Modelling buildings at LoD2 with the AHN point cloud is also possible, but results in models with slightly less details as some roof faces are no longer recognised. Many professional users of geo-information require the 3D models to be consistent with already available 2D building outlines. To achieve this at a national level the large scale base data (BGT) with 10-20 cm accuracy will be appropriate. In the reported experiments we did, however, not make use of this data, but obtained the building outlines from the lidar points classified as building roof.

3.2. Building modeling using roof topology graphs

Earlier work on building modelling using roof topology graphs has been presented by Verma et al. (2006) and Oude Elberink (2010). The work demonstrated the robustness of the modelling approach, but also showed that the available shapes of building parts in the target library did not allow a modelling of buildings with some more complex shapes. To increase the flexibility of modelling we

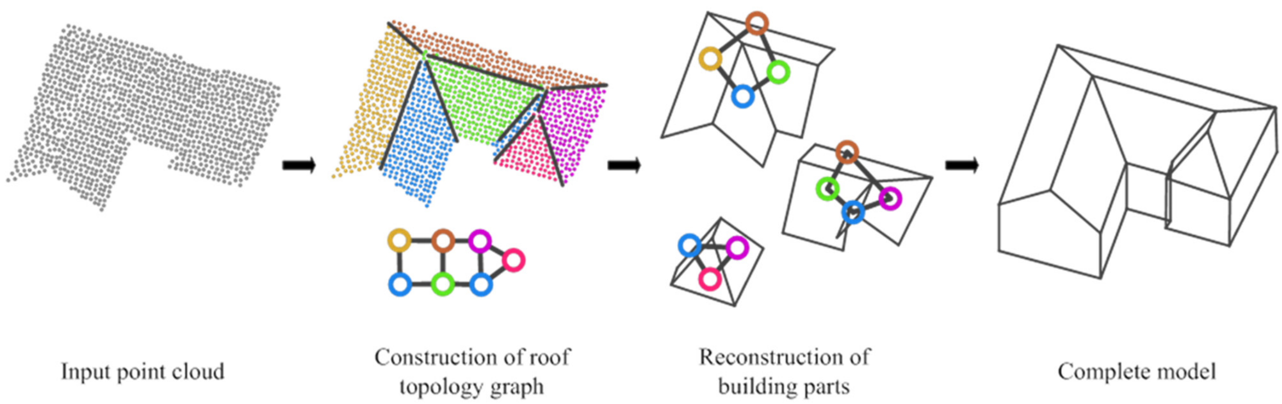

Figure 5: The topology of segmented roof point clouds serves to recognise shapes of roof parts.

proposed to include the simplest possible shapes into the target library (Xiong et al., 2014). In terms of the roof topology graph this applies that only minimal cycles of the graph and loose edges and nodes are stored as targets. Similar ideas have been developed in parallel by Perera and Maas (2014). The building modelling procedure is visualized in Figure 5. A point cloud of a roof is segmented into planar faces. Intersection lines are detected and lead to the roof topology graph describing the adjacency relationships between the roof faces. This graph is decomposed into minimal cycles that all correspond to a simple roof part. The combination of the recognised target shapes results in the complete 3D building model.

3.3. Automated understanding and correction of modelling errors

The success rate of the use of target shapes to recognise building parts depends on the correctness of the roof topology graphs that are extracted from the segmented point cloud. Because of the high pulse frequencies of modern laser scanners the point densities are typically high and allow a very reliable extraction of roof faces. Yet, in some cases roofs have very small roof faces that are not sampled dense enough to allow their detection in the point cloud. This will then result in missing nodes in the roof topology graph as well as missing and/or incorrect edges (adjacency relationships). Another failure to recognise the correct topology is shown in the left picture of Figure 6 where an intersection line is incorrectly hypothesised between two roof faces belonging to dormers on opposite sides of the gable roof. This results in an additional edge in the roof topology graph and leads to an incorrect model (2nd picture). When inspecting the errors made in the automated modelling using roof topology graphs we noticed that several types of errors were repeated in different buildings. Instead of asking an operator to correct the same type of error over and over again, we designed an algorithm that is capable of recognising a common type of error and correcting it (Xiong et al., 2014). For this purpose an operator corrects one instance of each type of error and stores the corrected roof topology graph together with the erroneous one in a database. The automated correction procedure starts with a characterisation of the building model quality by analysing the height differences between the building model and the point cloud. This results in pattern of well and less well fitting roof faces and intersection edges. If some parts of a building model are wrong, this pattern is compared against those of the errors stored in the database. When an error type is recognised, the corresponding corrected version is taken from the database to revise the roof topology graph of the building at hand. Afterwards the building model quality is assessed again to determine whether the model improved. In the example of Figure 6 the edge in the roof topology graph corresponding to the incorrectly detected intersection line is removed and the building model can be reconstructed correctly.

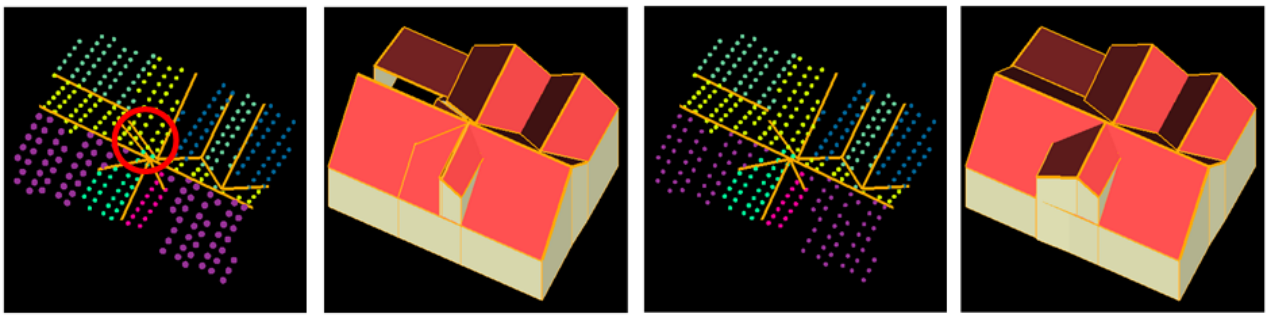

Figure 6: Correction of detected roof topology. Left: incorrect detected intersections of roof faces and the corresponding wrong model. Right: corrected intersections and model.

3.4. Interactive editing tools

In many cases where the automated methods for model reconstruction and error correction fail, the reconstructed topology graphs are still largely correct. Instead of modelling an entire building by hand, it is far more efficient to interactively correct the roof topology graph and then automatically reconstruct the building model. For this purpose we developed interactive editing tools that allow deletion and insertion of nodes and edges in a roof topology graph with only a few mouse clicks. The whole building modelling strategy then consists of three steps: (1) a completed automated building reconstruction based on decomposition of roof topology graphs in minimal cycles, (2) automated analysis of building model quality and recognition and correction of frequently occurring common types of errors, and (3) interactive editing of roof topology graphs for building models that still do not pass the quality test.

3.5. Results

The above sketched procedure for 3D building modelling has been applied to 9366 building objects in Enschede, the Netherlands, about one quarter of the total number of buildings in the city. To assess the quality of the automatic reconstruction, the building models are compared to the point cloud. A building model is considered to have an acceptable quality if there is no area larger than 3 m^2 with points that deviate more than 0.3 m from the model. Deviations caused by small dormers, chimneys and antennas are therefore not considered as errors. The automated modelling with the minimal cycles in roof topology graphs resulted in 93% correctly modelled buildings according to the above criterion. Application of the method for the automatic detection and correction of errors increased the acceptance rate to 95%. Thus, only 5% of the building models required interaction with an operator. The major reasons for failure of the automatic reconstruction were the lack of data on some surfaces with materials that largely absorb the laser pulses as well as the incorrect or missing detection of multiple small adjacent roof surfaces. In case of a single small roof surface, the automatic correction method may be able to successfully detect the error and repair the building model. However, when multiple roof surfaces (i.e. nodes in the roof topology graph) are missing, the operator's assistance is required. An example of an area with LoD2 building models is shown in Figure 7. Based on this experiment is was estimated that the modelling of all buildings in Enschede (158.000 inhabitants) would take five working days; half a day for the automatic reconstruction and error correction and the remaining time for editing the 5% of the buildings that could not be reconstructed automatically.

Figure 7: LoD2 buildings of a part of Enschede.

4. CONCLUDING REMARKS

This paper described the approach taken in the Netherlands to obtain a national 3D landscape model by fusing the national elevation data with the national 2D topographic database. Although the computational effort is large, it proved to be feasible. Several aspects are under discussion for further work.

While the LoD1 models are appropriate at the 1:10,000 scale for the fast majority of the buildings, some larger landmark buildings often consist of parts with very different heights. Modelling those buildings with a single flat roof results in models that are visually unattractive. For those building models it may be more appropriate to split the 2D building outline into multiple partitions corresponding to building parts with different heights prior to the 3D modelling.

For a next version of the national landscape model it needs to be decided whether the procedure should be based on change detection and selectively updating the changed areas or whether the landscape model should simply be regenerated for the whole country. The latter approach is easier to implement, but will lead to slightly changed building models for unchanged buildings when new elevation data becomes available.

The current modelling made use of the AHN-2 dataset obtained by airborne laser scanning. This elevation data currently has an update cycle of 5 years. It will be investigated whether the use of point clouds obtained by dense matching of annually acquired aerial photographs will have a sufficient quality for updating the national landscape model more frequently.

Furthermore, the quality of the 3D models needs to be inspected further and improved. This contains the validation of the consistency of the 3D models (Ohori et al., 2012) and the more accurate modelling of object surfaces (e.g. in case of outliers in vegetation and ground surfaces).

For the LoD2 building models we now only made use of high point density point clouds. However, when working with 3D geo-information, users typically require the 3D building models to be consistent with the already available 2D building outlines. This is still a challenging issue to be solved as small errors in the building outlines or in the roof faces obtained from the point cloud will then result in 3D building models with rather uncommon and visually unattractive topologies of roof faces.

At the current level of automation scaling the LoD2 modelling up to the national level is still another challenge. In terms of population Enschede is about 1% of the Netherlands. Assuming that Enschede has an average building complexity and an average ratio of inhabitants and buildings, the current modelling strategy would require around 450 days of interactive correction of building models to make LoD2 models available for the whole of the Netherlands. Although this would be feasible, an increase of the automatic modelling rate from the current 95% to e.g. 98% would make such a modelling project financially more attractive.

5. REFERENCES

Geonovum (2011): CityGML is the standard for 3D geo-information in the Netherlands. http://www.geonovum.nl/sites/default/files/3d_pilot_artikel_engels_2.pdf (Accessed 15 June, 2015).

Haala, N. and Kada, M. (2010): An update on automatic 3D building reconstruction. ISPRS Journal of Photogrammetry and Remote Sensing, 65, pp. 570-580.

Ohori, K. A., Ledoux, H. and Meijers, M. (2012): Validation and automatic repair of planar partitions using a constrained triangulation. Photogrammetrie, Fernerkundung, Geoinformation, 1 (5), pp. 613-630.

Oude Elberink, S. J. (2010) Acquisition of 3D topography: automated 3D road and building reconstruction using airborne laser scanner data and topographic maps. PhD thesis, University of Twente, Faculty of Geo-Information and Earth Observation (ITC).

Oude Elberink, S. J., Stoter, J. E., Ledoux, H. and Commandeur, T. (2013): Generation and dissemination of a national virtual 3D city and landscape model for the Netherlands. Photogrammetric Engineering & Remote Sensing (PE&RS), 79 (2), pp. 147-158.

Oude Elberink, S. J. and Vosselman, G. (2009): 3D information extraction from laser point clouds covering complex road junctions. The Photogrammetric Record, 24 (125), pp. 23-36.

Perera, G. S. N. and Maas, H.-G. (2014): Cycle graph analysis for 3D roof structure modelling: Concepts and performance. ISPRS Journal of Photogrammetry and Remote Sensing, 93, pp. 213-226.

Stoter, J., Ledoux, H., Reuvers, M., van den Brink, L., Klooster, R., Janssen, P., Beetz, J., Penninga, F., Vosselman, G. (2013): Establishing and implementing a national 3D standard in The Netherlands. Photogrammetrie – Fernerkundung – Geoinformation (PFG), 2013 (4), pp. 381-392.

Stoter, J. E., Vosselman, G., Dahmen, C., Oude Elberink, S.J. and Ledoux, H. (2014): CityGML Implementation Specifications for a Countrywide 3D Data Set. Photogrammetric Engineering & Remote Sensing 80, (11), pp. 1069-1077.

Verma, V., Kumar, R., Hsu, S. (2006): 3D building detection and modeling from aerial LiDAR data. Proceedings of the 2006 IEEE Computer Society Conference on Computer Vision and Pattern Recognition, CVPR'06, IEEE Computer Society, Washington, DC, pp. 2213-2220.

Vosselman, G. (2012): Automated planimetric quality control in high accuracy airborne laser scanning surveys. ISPRS Journal of Photogrammetry and Remote Sensing, 74, pp. 90-100.

Xiong, B., Oude Elberink, S.J. and Vosselman, G. (2014): A graph edit dictionary for correcting errors in roof topology graphs reconstructed from point clouds. ISPRS Journal of Photogrammetry and Remote Sensing, 93, pp. 227-242.