DETECTION OF BUILDINGS AT AIRPORT SITES USING IMAGES & LIDAR DATA AND A COMBINATION OF VARIOUS METHODS

Demir, N.[1], Poli, D.[2], Baltsavias, E.[1]
1- (nusret,manos@geod.baug.ethz.ch)
Institute of Geodesy and Photogrammetry, ETH Zurich, CH-8093, Zurich, Switzerland
2- (daniela.poli@jrc.ec.europa.eu)
European Commission - Joint Research Center, Ispra (VA), Italy

KEY WORDS: DTMs/DSMs, Lidar Data Processing, Multispectral Classification, Image Matching, Information Fusion, Object Detection, Buildings

ABSTRACT:

In this work, we focus on the detection of buildings, by combining information from aerial images and Lidar data. We applied four different methods on a dataset located at Zurich Airport, Switzerland. The first method is based on DSM/DTM comparison in combination with NDVI analysis (Method 1). The second one is a supervised multispectral classification refined with a normalized DSM (Method 2). The third approach uses voids in Lidar DTM and NDVI classification (Method 3), while the last method is based on the analysis of the density of the raw Lidar DTM and DSM data (Method 4). An improvement has been achieved by fusing the results of the different methods, taking into account their advantages and disadvantages. Edge information from images has also been used for quality improvement of the detected buildings. The accuracy of the building detection was evaluated by comparing the results with reference data, resulting in 94% detection and 7% omission errors for the building area.

1. INTRODUCTION

In this work, we focus on the building detection for airport sites. The acquisition of a reliable geospatial reference database of airports, and in particular the automatic extraction of buildings and obstacles at airports, both have a critical role for aviation safety. Often, 3D information of airports is not available, not accurate enough, not complete, or not updated. Thus, methods are needed for generation of accurate and complete 3D geodata with high degree of automation. In particular, buildings and trees are considered as obstacles, so they should be correctly extracted. In this work, we focus on the detection of buildings, as a first step for their 3D extraction. There are several methods applied for this purpose, based on image and/or airborne Lidar data. In our approach, buildings are detected in aerial images and Lidar data through multiple methods using multispectral image classification, DSM (Digital Surface Model) and DTM (Digital Terrain Model) comparisons and density analysis of the raw Lidar point cloud. The detection quality is improved by a combination of the results of the individual methods. This paper will give a brief overview of the related work on this subject. Then, after the description of the test area at Zurich Airport, Switzerland, the strategy and methodology will be presented and the results will be reported, compared and commented. This work is a part of the EU 6th Framework project PEGASE (Pegase, 2009).

2. PREVIOUS WORK

Aerial images and Lidar data are common sources for object extraction. In digital photogrammetry, features of objects are extracted using 3D information from image matching or DSM/DTM data, spectral, textural and other information sources. Pixel-based classification methods, either supervised or unsupervised, are mostly used for land-cover and man-made structure detections. For the classical methods e.g. minimum-distance, parallelepiped and maximum likelihood, detailed information can be found in (Lillesand and Kiefer, 1994).

In general, the major difficulty in using aerial images is the complexity and variability of objects and their form, especially in suburban and densely populated urban regions (Weidner and Foerstner, 1995).

Regarding Lidar, building and tree extraction is basically a filtering problem in the DSM (raw or interpolated) data. Some algorithms use raw data (Sohn and Dowman, 2002; Roggero, 2001; Axelsson, 2001; Vosselman and Maas, 2001; Sithole, 2001; Pfeifer et al., 1998), while others use interpolated data (Elmqvist et al., 2001; Brovelli et al., 2002; Wack and Wimmer, 2002). The use of raw or interpolated data can influence the performance of the filtering. The algorithms differ also in the number of points they use at a time. In addition, every filter makes an assumption about the structure of bare-earth points in a local neighbourhood. This assumption forms the concept of the filter (Sithole and Vosselman, 2003). The region-based methods use mostly segmentation techniques, like in Brovelli et al. (2002), or using Hough transformation (Tarsha-Kurdi et al., 2007). Some researchers use 2D maps as prior information for building extraction (Brenner, 2000; Haala and Brenner., 1999; Durupt and Taillandier, 2006; Schwalbe et al., 2005). Topographic maps provide outlines, classified polygons and topologic and 2D semantic information (Elberink and Vosselman, 2006).

In general, in order to overcome the limitations of image-based and Lidar-based techniques, it is of advantage to use a combination of these techniques. Sohn and Dowman (2007) used IKONOS images to find building regions before extracting them from Lidar data. Straub (2004) combines information from infrared imagery and Lidar data to extract trees. Rottensteiner et al. (2005) evaluate a method for building detection by the Dempster-Shafer fusion of Lidar data and multispectral images. They improved the overall correctness of the results by fusing Lidar data with multispectral images.

Few commercial software packages allow automatic terrain, tree and building extraction from Lidar data. In TerraSCAN, a TIN is generated and progressively densified, the extraction of off-terrain points is performed using the angles between points to make the TIN facets and the other parameter is the distance to nearby facet nodes (Axelsson, 2001). In SCOP++, robust methods operate on the original data points and allow the simultaneous elimination of off-terrain points and terrain surface modelling (Kraus and Pfeifer, 1998).

In summary, most approaches try to find objects using single methods. In our strategy, this study suggests complying different methods using all available data with the focus on improving the results of one method by exploiting the results from the remaining ones.

3. INPUT DATA AND PREPROCESSING

The methods presented in this paper have been tested on a dataset of the Zurich airport. The available data for this region are: 3D vector data of airport objects, colour and CIR (Colour InfraRed) images, Lidar DSM/DTM data (raw and grid interpolated). The characteristics of the input data can be seen in Table 1.

Image Data	**RGB**	**CIR**
Provider	Swissphoto	Swissphoto
Scale	1: 10‘000	1: 6‘000
Scan Resolution	14.5 microns	14.5 microns
Acquisition Date	July 2002	July 2002
Ground Sampling Distance (GSD) (cm)	14.5 cm	8.7 cm
Lidar Data	**DSM**	**DTM**
Provider	Swisstopo	Swisstopo
Type	Raw & grid	Raw & grid
Raw point density & Grid Spacing	1 pt / 2 sqm & 2m	1 pt / 2 sqm & 2m
Acquisition Date	Feb. 2002	Feb. 2002
Vector data	Only for validation purposes	
Provider	Unique Co.	
Horizontal / Vertical Accuracy (2 sigma)	20 / 25 cm	

Table 1. Input data characteristics.

The 3D vector data describe buildings (including airport parking buildings and airport trestlework structures). It has been produced from stereo aerial images using the semi-automatic approach with the CC-Modeler software (Gruen and Wang, 1998). Some additional reference buildings outside the airport perimeter were collected using CIR images with stereo measurement by using LPS software. The images have been firstly radiometrically preprocessed (noise reduction and contrast enhancement), then the DSM was generated with the software package SAT-PP, developed at the Institute of Geodesy and Photogrammetry, ETH Zurich (Zhang, 2005). For the selection of the optimum band for matching, we considered the GSD, and the quality of each spectral channel based on visual checking and histogram statistics. Finally, the NIR band was selected for DSM generation. The final DSM was generated with 50cm grid spacing. Using this DSM, CIR orthoimages were produced with 12.5cm ground sampling distance. Lidar raw data (DTM and DSM) have been acquired with “leaves off”. The DSM point cloud includes all Lidar points (including points on terrain, tree branches etc.). The DTM data includes only points on the ground, so it has holes at building positions and less density at tree positions. The height accuracy (one standard deviation) is 0.5 m generally, and 1.5 m at trees and buildings, the latter referring only to the DSM. The grid DSM and DTM were interpolated from the original raw data by Swisstopo with the Terrascan commercial software.

4. BUILDING DETECTION

Four different approaches have been applied to exploit the information contained in the image and Lidar data, extract different objects and finally buildings. The first method is based on DSM/DTM comparison in combination with NDVI (Normalised Difference Vegetation Index) analysis for building detection. The second approach is a supervised multispectral classification refined with height information from Lidar data and image-based DSM. The third method uses voids in Lidar DTM and NDVI classification. The last method is based on the analysis of the density of the raw DSM Lidar data. The accuracy of the building detection process was evaluated by comparing the results with the reference data and computing the percentage of data correctly extracted and the percentage of reference data not extracted.

4.1 DSM/DTM and NDVI (Method 1)

By subtracting the DTM from the DSM, a so-called normalized DSM (nDSM) is generated, which describes the above-ground objects, including buildings and trees. As DSM, the surface model generated by SAT-PP and as DTM the Lidar DTM grid were used. NDVI image has been generated using the NIR and R bands. A standard unsupervised (ISODATA) classification was used to extract vegetation from NDVI image. The intersection of the nDSM with NDVI should correspond to trees. By subtracting the resulting trees from the nDSM, the buildings are obtained. 83% of building class pixels were correctly classified, while all of 109 buildings have been detected but not fully, the omission error is 7% . Within the detected buildings, some other objects, such as aircrafts and vehicles, were included. The extracted buildings are shown in Figure 1.

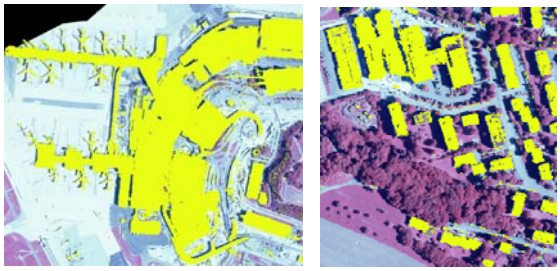

Figure 1. Building detection result from method 1. (Left: airport buildings, Right: residential area).

4.2 Supervised classification and use of nDSM (Method 2)

The basic idea of this method is to combine the results from a supervised classification with the height information contained in the nDSM. Supervised classification methods are preferable to unsupervised ones, because the target of the project is to detect well-defined standard target classes (airport buildings, airport corridors, bare ground, grass, trees, roads, residential houses, shadows etc.), present at airport sites. The training areas were selected manually using AOI (Area Of Interest) tools within the ERDAS Imagine commercial software (Kloer, 1994). Among the available image bands for classification (R, G and B from colour images and NIR, R and G bands from CIR images), only the bands from CIR images were used due to their better resolution and the presence of NIR channel (indispensable for

vegetation detection). In addition, new synthetic bands were generated from the selected channels: a) 3 images from principal component analysis (PC1, PC2, PC3); b) one image from NDVI computation using the NIR-R channels and c) one saturation image (S) obtained by converting the NIR-R-G channels in the IHS (Intensity, Hue, Saturation) colour space. The separability of the target classes was analyzed through use of plots by mean and standard deviation for each class and channel and divergence matrix analysis of all possible combinations of the three CIR channels and the additional channels, mentioned above. The analysis showed that:

- G and PC2 have high correlation with other bands
- NIR-R-PC1 is the best combination based on the plot analysis
- NIR band shows good separability based on the divergence analysis
- PC1-NDVI-S combination shows best separability over three-band combinations based on the divergence analysis.

Therefore, the combination NIR-R-PC1-NDVI -S was selected for classification. The maximum likelihood classification method was used. As expected from their low values in the divergence matrix, grass and trees, airport buildings and residential houses, airport corridors and bare ground, airport buildings and bare ground could not be separated. Using the height information from nDSM, airport ground and bare ground and roads were fused into "ground" and airport buildings with residential houses into "buildings", while trees and grass, as well as buildings and ground could be separated. The final classification is shown in Figure 2. 84% of the building class is correctly classified, while All of 109 buildings have been detected but not fully, the omission error is 9% . Aircrafts and vehicles are again mixed with buildings.

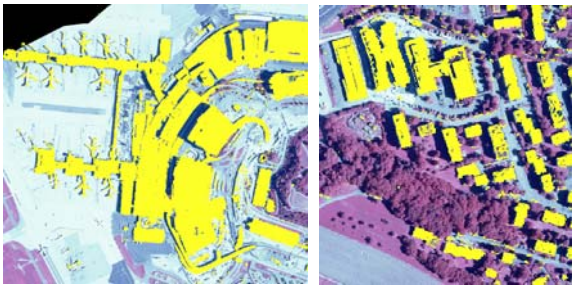

Figure 2. Building detection result from method 2. (Left: airport buildings, Right: residential area).

4.3 Building detection using density of raw Lidar DTM and NDVI (Method 3)

Buildings and other objects, like high or dense trees, vehicles, aircrafts, etc. are characterized by null or very low density in the DTM point cloud. Using the vegetation class from NDVI channel as a mask, the areas covered by trees are eliminated, while small objects (aircrafts, vehicles) are eliminated by deleting them, if their area is smaller than $25m^2$. Thus, only buildings remain (Figure 3). 85% of building class pixels are correctly classified, while 108 of 109 buildings have been detected but not fully extracted, the omission error is 8% .

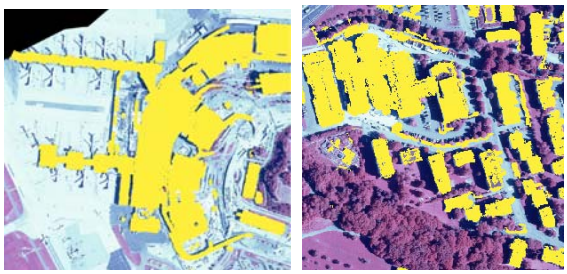

Figure 3. Building detection result from method 3. (Left: airport buildings, Right: residential area).

4.4 Building and tree detection from Lidar data (Method 4)

As mentioned above, in the raw DSM data the point density is generally much higher at trees than at open terrain or buildings. On the other hand, tree areas have low horizontal point density in the raw DTM data. We start from regions that are voids or have low density in the raw DTM (see Method 3). These regions represent mainly buildings and trees and are used as mask to select the raw DSM points for further analysis. In the next step, we used a search window over the raw Lidar DSM data with a size of 5 m x 5 m. Neighboring windows have an overlap of 50%. The window size has a relation with the number of points in the window and the number of the points in the search window affects the quality of the detection result. The method uses all points in the window and labels them as tree if all parameters below have been met. The size of $25m^2$ has been agreed to be enough to extract one single tree. A bigger size may result in wrong detection especially in areas where the buildings are neighboring with single trees. On the other hand, the data has low point density: 1 pt / 2 m^2, that means about 13 pts / 25 m^2. A smaller size will contain less points and this may not be enough for the detection.

The points in each search window are projected onto the xz and yz planes and divided for each projection in eight equal sub-regions using x_{min}, x_{mid}, x_{max}, z_{min} z_{mid1} z_{mid2} z_{mid3} z_{max} as boundary values of sub-regions, with $x_{mid} = x_{min} + 2.5m$, $x_{max} = x_{mid} + 2.5m$, $z_{mid1}=z_{min}+(z_{max}-z_{min})/4$, $z_{mid2} =z_{min}+2*(z_{max}-z_{min})/4$, $z_{mid3}=z_{min}+3*(z_{max}-z_{min})/4$ and similarly for the yz projection. The density in the eight sub-regions is computed. The first step is the detection of trees and the second the subtraction of tree points from all off-terrain points. The trees have been extracted by four different parameters. The parameters have been calculated using tree-masked areas of the raw Lidar DSM data. The tree mask has been generated by Method 2. Then, the calculated parameters (the average of all search windows) have been applied to the raw Lidar DSM data for detection of trees.

The first parameter (s) is similarity of surface normal vectors. We assume that the tree points would not fit to a plane. With selection of three random points in the search window, the surface normal vectors have been calculated n (number of points in search window) times. Then, all calculated vectors have been compared among each other. In case of similar value of compared vectors, the similarity value was increased by adding 1. In the tree masked points, the parameter (s) has been calculated as smaller than 2. The second parameter (vd) is the number of the eight sub-regions which contain at least one point. The trees have high Lidar point density vertically. Thus, at trees more sub-regions contain Lidar points. Using the tree mask, we have observed that at least 5 out of the 8 sub-regions contain points. Thus, the parameter (vd) has been selected as

vd>4. The third parameter (z) is the tree height. Using the tree mask from multispectral classification, we calculated the minimum tree height as 3m. The fourth parameter (d) is the point density. The minimum point density has been calculated for the tree masked areas as 20points/ 25m^2. By applying these four parameters to the raw DSM Lidar data, the tree points have been extracted and eliminated from all off-terrain points to extract the buildings. The workflow can be seen in Figure 4.

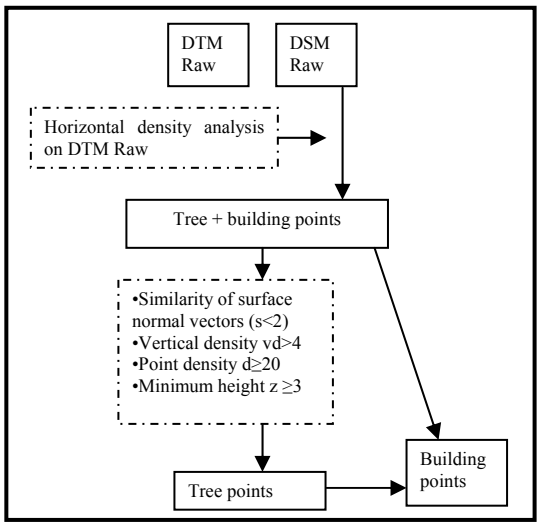

Figure 4. Workflow of detection of buildings in method 4

The density of point cloud directly affects the quality of the result. In addition, some tree areas could not be extracted because of the low point density of the Lidar data. The accuracy analysis shows that 84% of buildings area are correctly extracted, while 100 of 109 buildings have been detected but not fully extracted, the omission error is 17% .(Figure 5).

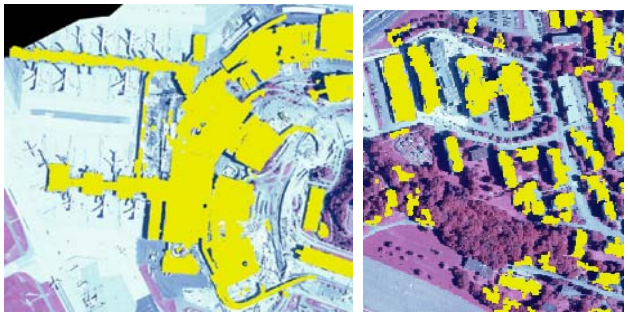

Figure 5. Building detection result from method 4. (Left: airport buildings, Right: residential area).

5. ANALYSIS OF RESULTS

Each method shows similar performance with differences in completeness. The reasons of the failures for correctness and completeness of each method can be seen in Table 2. The improvement of the results is performed by taking into account the advantages and disadvantages of the methods.

	Correctness Failure Reasons	Completeness Failure Reasons
M1	Airplanes/Other moving objects /shadow on vegetation/construction process	Vegetation on roofs, lack of some parts of buildings which are being constructed.
M2	Airplanes/Other moving objects/construction process	Vegetation on roofs, shadow on roofs, lack of some parts of buildings being constructed.
M3	Moving objects (esp. car series in parking lots)/ other man-made structures (highways etc.)	Vegetation on roofs, temporal difference with reference data
M4	Tree groups which could not be extracted and eliminated	Non-detection of small buildings (problem related to low point density), detection of walls as vegetation, temporal difference with reference data

Table 2. The reasons of the failures regarding correctness and completeness for each method (M: Method).

Regarding completeness, the reference data has been generated using aerial images, and some buildings are in construction process. Reference data has been provided from Unique Company and they have produced it using aerial images. But, in the construction areas, these buildings were measured as fully completed, although they were only partly constructed in reality. This increases the omission error especially for the results of the methods 1 and 2 which use aerial images. On the other hand, due to the temporal difference between the reference vector and Lidar data, the completeness of Lidar-based methods (methods 3 and 4) has also been negatively affected.

5.1. Combination of the methods

The results from each method have been combined according to their failures for different types of objects. Intersection of all methods gives the best correctness, while the union of the methods gives the best completeness. The combination of the results has been performed for achieving the best correctness with the best completeness.

(1∩2): While method 2 does not include the errors resulted by the shadow on vegetation, the intersection of these two methods eliminates the problem of shadow on-vegetation (in Figure 12, R1). The correctness of extracted buildings from this combination is 86%, and the omission error is 12%.

(1∩2) ∩4: This combination eliminates the airplane objects from the detection result (Figure 6). Consequently, another advantage of this combination is that it reduces the omission errors which arise from the construction process on some buildings, i.e. multitemporal differences. The correctness of extracted buildings from this result is 96%, and the omission error is 20% (in Figure 12, R2).

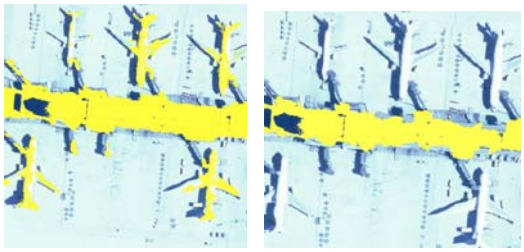

Figure 6. Left: Airplanes which were detected as buildings in (1∩2), Right: Elimination of airplanes with (1∩2) ∩4.

((1∩2) ∩4) ∪ 3: Shadow regions on buildings are replaced with building regions and by this combination (Figure 7). Since method 3 brings the buildings which could not be detected well by method 4, and method 3 is not influenced by shadow, this combination provides better completeness (in Figure 12, R3).

Figure 7. Left: buildings without the regions which covered by shadow in ((1∩2) ∩4), Right: more complete roofs with ((1∩2) ∩4) U 3.

After the union process with the results of the Method-3, the vegetation on the roof tops is still a problem. Intersection of the nDSM and the NDVI algorithms provides the tree and vegetation regions on the roof tops. Intersection of the extracted vegetated regions with building polygons of the Method-4 results in the roof regions which contain vegetation (Figure 8).

Figure 8. Roof regions which contain vegetation.

After adding the roof regions which contain vegetation into the detection result (in Figure 12, R4), the correctness and completeness values are 85% and 7%. As mentioned before, since method 2 have detected all buildings although not fully, the final building polygons should overlap the results from method 2. If the building polygons of result (R4) do not overlap with the results of method 2, they are eliminated. The correctness of the results is improved to 91% and the omission is 7% (Figure 12, R5).

5.2. Using edge information for improvement of correctness

Image data provide edge information, and this can be used to find the precise outlines of the buildings. Firstly, the Canny edge detector (Canny, 1986) has been applied on the orthoimages. The edges have been split into straight lines using corner points which were detected by corner detection (Harris and Stephens, 1988). This has been performed using the Gandalf image processing library (Gandalf, 2009). The straight lines which are smaller than 1 m. have been considered as noise and they have been deleted. The straight lines which may belong to building outlines have been selected using the outline of the detection result (which comes from the combination of methods) and a 2m buffer zone (1m inside, 1 m outside of the building outline). If the straight lines are neighbours in the buffer zone, the longest straight line has been selected. There is an exception for this neighboring criterion: the start or end point of a straight line should not be the closest point to the neighbouring line. With this exception, we avoid the elimination of lines, which are almost collinear (Figure 9).

Figure 9. Left: the straight lines which may belong to the building outline (yellow) and Right: long lines (red).

After selection of the straight lines, they have been converted to closed polygons. For the conversion to polygons, a sorting of the lines in clock-wise direction is used. To perform sorting, the travelling salesman convex hull algorithm (Deineko et al., 1992) has been applied. After closing the polygons, we separate the lines into those that were detected from the images (red) and the ones added by this algorithm (blue) (see Figure 10). The red straight lines, which are shorter than 10 m. and form an acute angle (between 1 and 80 degrees), are eliminated (Figure 10), as well as all blue lines.

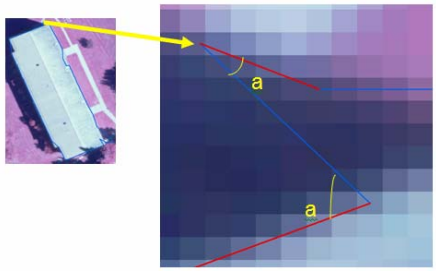

Figure 10. Line elimination procedure when the line length is shorter than 10 meters and has acute angle with its neighbouring lines (red: eliminated lines, blue: lines added by the travelling salesman algorithm, yellow: acute angle).

If two red lines form an acute angle and are shorter than 10 m., then both lines are eliminated. After this elimination, the travelling salesman convex hull algorithm has been applied again using the non-eliminated red lines and generated the refined building polygons (Figure 11).

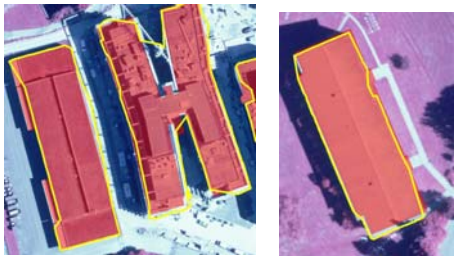

Figure 11. Final building polygons (yellow)., and reference data (red).

After this process, the correctness has been improved to 94% with remaining 7% omission error (Figure 12, R6). However, it has not been applied on all 109 buildings of this test yet, due to time restrictions, while it has shortcomings, as the travelling salesman algorithm does not use any input data information for forming closed polygons.

5.3 Final results

The rule-based system for the combination of methods can be seen in Figure 12.

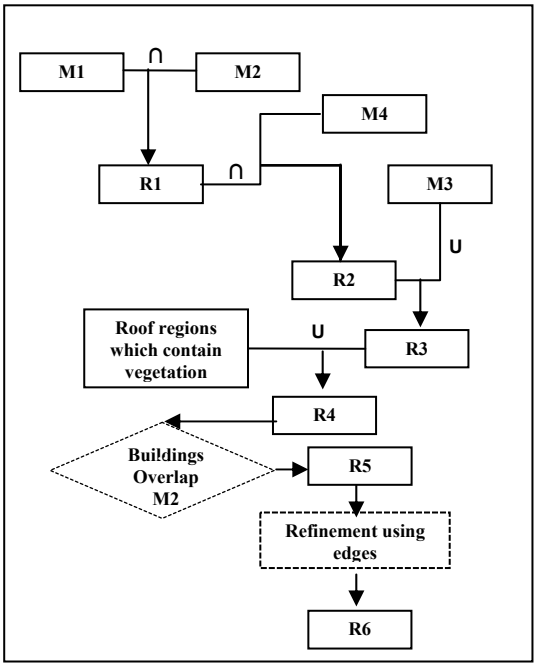

Figure 12. Combination of the methods. R: result from combination, M: Method**.**

Table 3 gives a summary of the correctness and omission percentages of the various detection methods.

	Correctness (%)	Omission (%)		Correctness (%)	Omission (%)
M1	83	7	R1	86	12
M2	84	9	R2	96	20
M3	85	8	R3	85	8
M4	84	17	R4	85	7
			R5	91	7
			R6	94	7

Table 3. Summary of the correctness and omission percentages.

6. CONCLUSIONS

In this paper, different methods for object detection (mainly buildings) in Lidar data and aerial images have been presented. In each method, the basic idea was to get first preliminary results and improve them later using the results of the other methods. The methods have been tested on a dataset located at Zurich Airport, Switzerland, containing RGB and CIR, Lidar DTM and DSM point clouds and regular grids and building vector data for accuracy assessment. The results from each method have been combined according to their error characteristics. Edges have been used for further improvement of the detected building outlines. Finally, the correctness of detection has been 94% with remaining 7% omission error that mostly comes from construction process on airport buildings. Future work will focus on the improvement of use of edges, using the Lidar DSM to eliminate lines which don't belong to buildings and 3D building roof modeling.

ACKNOWLEDGEMENTS

This work has been supported by the EU FP6 project Pegase. We acknowledge data provided by Swisstopo and Unique Company (Airport Zurich).

REFERENCES

Axelsson, P., 2001. Ground estimation of laser data using adaptive TIN-models. Proc. of OEEPE workshop on airborne laserscanning and interferometric SAR for detailed digital elevation models, 1-3 March, Stockholm, Sweden, pp. 185-208.

Brenner, C., 2000. Towards fully automatic generation of city models. Int. Archives of Photogrammetry and Remote Sensing, Vol 33, Part B3/1, pp.85-92.

Brovelli, M.A., Cannata, M., Longoni U.M., 2002. Managing and processing Lidar data within GRASS. Proc. of the GRASS Users Conf. 2002, Trento, Italy. http://citeseer.ist.psu.edu/541369.html (accessed 02 July 2009).

Canny, J., 1986. A computational approach to edge detection. IEEE Trans. Pattern Anal. Machine Intell., 8(6), 679–698.

Deineko, V., Van Dal, R., Rote, G., 1992. The convex-hull-and-line traveling salesman problem: A solvable case. Information Processing Letters, 51 (3), 141–148. http://citeseer.ist.psu.edu/old/286488.html (accessed 02 July 2009).

Durupt, M., Taillandier, F., 2006. Automatic building reconstruction from a digital elevation model and cadastral data : An operational approach. IAPRS , Vol. 36, Part 3, pp. 142-147. http://www.isprs.org/commission3/proceedings06/singlepapers/O_14.pdf (accessed 02 July 2009).

Elberink, S.,O., Vosselman, G., 2006, 3D Modelling of Topographic Objects by Fusing 2D Maps and LIDAR Data, IAPRS* Vol. 36, Part 4, pp. 199-204. http://intranet.itc.nl/papers/2006/conf/vosselman_3D.pdf (accessed 02 July 2009).

Elmqvist, M., Jungrt, E., Lantz, F., Persson, A., Soderman, U., 2001. Terrain modelling and analysis using laser scanner data. IAPRS*, Vol. 34, Part 3/W4, pp. 219-227. http://www.isprs.org/commission3/annapolis/pdf/Elmqvist.pdf (accessed 02 July 2009).

Gandalf, 2009. http://gandalf-library.sourceforge.net/ (accessed 27 June 2009).

Gruen, A., Wang, X., 1998. CC-Modeler: A topology generator for 3-D city models. ISPRS Journal of Photogrammetry & Remote Sensing 53(5), 286-295. http://linkinghub.elsevier.com/retrieve/pii/S0924271698000112 (accessed 02 July 2009).

Haala, N., and Brenner, C., 1999. Virtual city models from Laser altimeter and 2D map data. Photogrammetric Engineering & Remote Sensing 65 (7), 787-795. http://www.ifp.uni-stuttgart.de/publications/1999/norbert_ohio.pdf (accessed 02 July 2009).