Building reconstruction and visualization from LIDAR data

G. Forlani [a,] C. Nardinocchi [b], M. Scaioni [c], P. Zingaretti [d]

[a] Dept. of Civil Engineering, University of Parma – gianfranco.forlani@unipr.it
[b] Dept. ITS, University of Rome "La Sapienza" – carla.nardinocchi@uniroma1.it
[c] Dept. IIAR, Politecnico of Milan – marco.scaioni@polimi.it
[d] Dept. of Informatic, University of Ancona – zinga@inform.unian.it

Commission V, WG V/4 & IC III/5

KEY WORDS: Range images, Segmentation, Terrain Classification, Building extraction and 3D reconstruction.

ABSTRACT:

A strategy for building extraction and building reconstruction from dense LIDAR data is presented. Roofs are modeled as plane surfaces, connected along ridges and bordered by the eaves lines. Initial segmentation of terrain and non-terrain features is performed on grid data; after segmentation by region growing and region topology, classification is achieved based on a hierarchical set of aggregation rules. Areas labelled as potential buildings are further segmented in plane surfaces (roofs slopes) based on gradient orientation and plane fitting by RANSAC; a similar procedure applies to edge pixels to extract eaves lines. The topology of the roof slopes and walls is reconstructed, deriving also roof ridges and roof corners. A 3D model of the building is thereafter obtained which can be imported in CAD and visualization environments, such as those provided by the VTP group. Results from a laser scanning survey with a ground resolution of 1 m are presented for a suburban area in Pavia.

1. INTRODUCTION

LIDAR is promoting itself as the most effective data acquisition technology for high accuracy, high resolution DTMs currently available. Short production times are obtained thanks to system and sensor integration in raw data acquisition and to the development of tools and algorithms to automate data processing. Interactive editing is nevertheless necessary because data classification often fails to filter out non-terrain points: there is a great interest in improving algorithm's performance on this topic. Besides, the dramatic amount of data to handle puts hardware to its limits and feeds demands for faster and flexible visualization and data manipulation techniques.
In this paper we will concentrate on data classification, where at least three major items (proper terrain, vegetation and buildings) should be identified, and on the production of 3D models of buildings from LIDAR data. Indeed, besides DTM production, the dense 3D point cloud acquired by LIDAR is seen as alternative to aerial images in building extraction and building reconstruction for the generation of 3D city models.
From urban planning to the set up of telecommunications networks, 3D representations of the urban environment (3D city models) are produced and input in spatial databases, visualization or virtual reality tools, wave propagation simulators, etc. Depending on the purpose, buildings can be represented by simple boxes, with no pictorial information; roof modeling, photorealistic texture or even rectified images may be required, at least in some areas, for virtual tours. This affects production costs as well as data acquisition and modelling techniques. If building elevations and digital ground plans are available, simple solid models can be easily produced; if a DTM is available, more realistic perspective views can be generated, possibly draping or mapping texture on the DTM (colour raster maps or orthoimages are a popular choice) to improve the visual representation. Roof modelling, adding texture to building facades and including vegetation dramatically increases the production costs, calling for as much automation as possible.
Currently, either fully automatic as well as interactive building extraction systems are being developed. Aerial images have been the main data source for such systems up today; even with multiple overlapping images, though, sophisticated signal and feature based methods still face problems in an urban environment. Most segmentation problems arising from shadows and occlusions can be considerably reduced using range data, such as acquired by airborne laser scanners. In this paper, we want to present a strategy for automatic data classification and building extraction from LIDAR data only. Our target is the separation of raw data in several classes and the geometric reconstruction of polyhedral models of buildings, achieved by identifying roof slopes and eaves lines.
The paper is organized as follows: in Section 1, we describe the data set used to validate our strategy (Section 3); Section 2 summarizes terrain classification, building localization and roof segmentation. The solutions to problems concerning the construction of the 3D vector models of the buildings are discussed in Section 3, while the last section is devoted to the adoption of an efficient, interactive data exploration tool before the conclusions are drawn.
A detailed review of the existing methods for terrain classification and building reconstruction is beyond the scope of this paper: only some basic reference will be given to underline overlaps and differences between our and others' work.
As a rule, building reconstruction is performed by a preliminary segmentation, followed by data classification and generation of building hypotheses (building localization); each hypothesis is later checked for further evidence, measuring its consistency with a building model and reconstructing its shape and location. LIDAR has been taken into consideration in building reconstruction since mid nineties; Haala & Brenner, 1997, Brenner , 2000 use ground plans to detect buildings and form roof slope hypotheses; roof modelling is achieved by segmenting a Digital Surface Model (DSM) in planar surfaces

by RANSAC and reconstructing roof topology by a rule based approach. Weidner et al.' 1995 detect buildings by analysing the blobs in a normalized DSM (i.e. a DSM obtained by subtracting the DTM from the original DSM) and separates buildings from vegetation looking at the variance of the local surface normals; building extraction is performed using parametric models (general polyhedral can be recovered) by extracting the building ground plan, estimating the height and selecting the appropriate model from a library by MDL principle.

2. DATA SET

Under a national research project aiming at the assessment of LIDAR technology for DTM production, height data were acquired in November 1999 over the city of Pavia with a Toposys laser scanning system; stereo aerial images were also gathered. The covered area is about 30 km^2 large and includes the old city centre, with narrow streets and very complex roof shapes, suburban areas with both high rise buildings and detached houses with trees; the Ticino river stretching in the southern part of the town and some countryside areas with farms.
The Toposys instrument flew at a height above ground of about 800 m in Est-West direction: with a scan angle of 14° the swath width is about 250 meters. This low scan angle has the advantage of a good penetration in narrow streets between buildings, which is typical in most of old Italian cities; several strips where flown with larger than normal overlaps and some in North-South direction, crossing all the others, with the purpose of studying strip registration. Because of the system characteristics, the pattern of the laser spots on a flat terrain would be regular, though anysotropic, with a spacing of about 1.6 m across track and of 0.15 m along track . The point density with these mission parameters is about 5 point/m^2.

3. BUILDING RECONSTRUCTION AND VISUALIZATION FROM LIDAR DATA

In the following, we present our strategy for data classification, building detection and building reconstruction based on LIDAR data only. Its central idea is to combine data classification and building extraction, putting the topological description of the data at the core of both tasks. At each stage of the segmentation process, the topological structure of the segmented regions is updated and relevant information encoded in a data base. Since, besides inner characteristics, neighbourhood relationships are crucial to separate and label regions either in building detection as well as in roof segmentation and roof modelling, we believe this is a key element to allow a rule-based scheme to succeed.
A second important aspect of our approach is that we use grid data to take advantage of their regularity in data processing, but we retain the possibility to go back to raw data when necessary, for example, to extract more accurate building contours. We chose as optimal a mesh size of 1x1 m; since several raw data points fall into a grid cell, the interpolation algorithm that is applied may privilege the lower points in DTM generation or the higher points in building extraction.
Each step of the procedure, namely terrain classification, building detection and building reconstruction, is now briefly addressed to give the overall picture; details will be given in the sequel.
Building reconstruction starts deriving the adjacency relationships between eaves and roof slopes; based on a set of rules, vertices of the roof are computed as the intersection of adjacent terns of planes (either sloping or vertical) and connected by edges leading to the wireframe representation of the roof and of the walls.

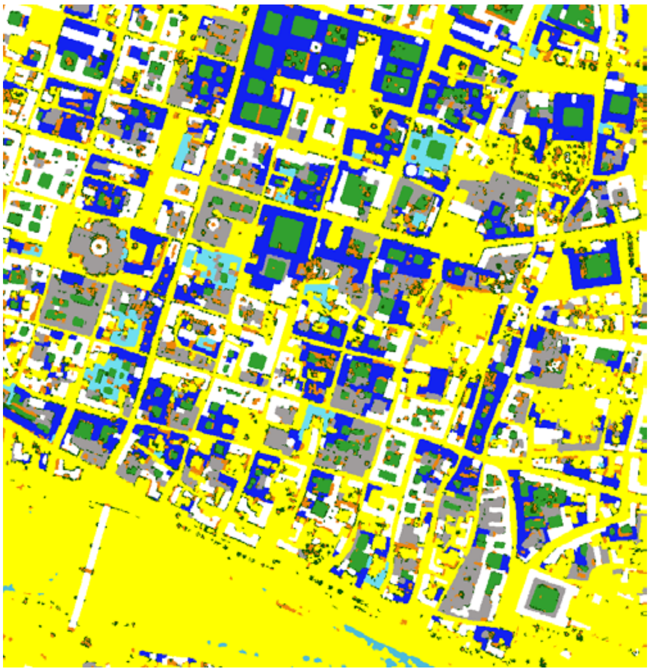

Fig. 1 – Results of data classification in the old city cente

3.1 Data segmentation and classification

As far as data classification is concerned, a primary segmentation is performed by aggregating pixels in smooth connected regions: nowhere in a region a pixel elevation differs from that of all neighbours by more than a threshold. Statistical parameters and shape descriptors of each region are computed and the relationships with adjacent regions recovered. The hierarchical application of a set of rules allows to assign regions to several classes: bare terrain, vegetation, buildings, courtyards, water, etc.
Elevated regions are labelled as possible buildings and further processed; since the building model assumes a building as a volume bounded by planes and raising with step edges above the nearby region, domes, some kind of industrial buildings and unconventional architectures will not be recognised.
We apply a region growing technique to the whole dataset, looking for connected sets (regions) of height data that are bounded by step edges. Typically, a threshold of 0.5 m is used to cluster the whole dataset.
Small regions, distinguished in noise, vegetation and elongated regions, are not further considered. The relationship between regions are then analysed to classify all the regions, according to a set of rules. These clusters are distinguished in terrain, courtyard and buildings. This does not end the detection stage: whether a region does actually represent a building or not is deferred to a later stage, when reconstruction is also under way.
Figure 1 shows the result of the classification in an urban-suburban area. The yellow pixels represent the terrain. Courtyards are in light green. White, dark blue, light blue and gray regions represent buildings with different heights. Vegetation is in dark green, while elongated regions, which may represent balconies, low walls, platform roofs or vegetation, are in orange.

3.2 Roof slope extraction

Elevated regions are segmented to find roof slopes: sub-regions with homogeneous gradient orientation are grouped. For each pixel of an elevated region the gradient orientation is computed. Then, to identify roof slopes, connected sets of pixels belonging to the same partition of the orientation space[2] are grouped into clusters. The cluster with less than three points are discarded.

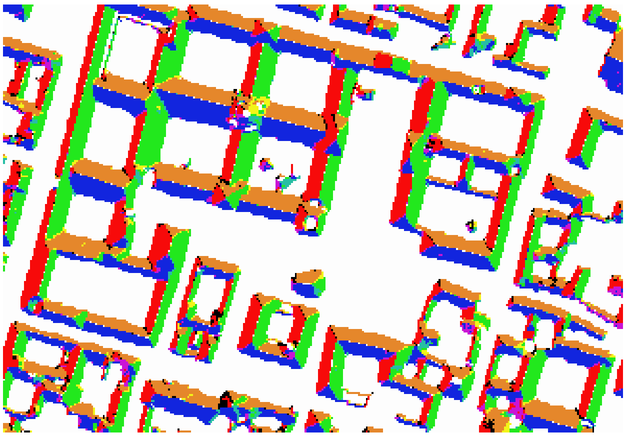

Figure 2 – Results of the roof slopes segmentation using a partition in 8 classes of the gradient orientation

Figure 2 shows the segmentation of the regions classified as elevated. Pitches with the same orientation are represented with the same color. If a region does not have at least a planar surface it is no more considered as a 'building'.

The points assigned to a cluster do not necessarily identify a single pitch. A cluster may indeed include several adjacent roof slopes, either slightly convergent (figure 4) or parallel (figure 3), because the slope may be different and the range of orientations within a cluster may be large. To discriminate such cases, we fit a plane to the cluster's points; if the residuals are larger than 15 cm (the accuracy of laser data in height), the RANSAC algorithm is applied. Figure 4 and 5 show two cases where the application of RANSAC corrects a wrong segmentation.

While this generally leads to good and robust identification of roof slopes, eaves are harder to derive by just grouping border pixels of elevates regions: the discrete sampling of the laser spots leads to poorly defined edges. Again we use RANSAC to group border pixels in line segments but we also check whether edge points belong to a single horizontal plane. If this is the case (as often it is), the building outlines will be computed by intersection of this plane with the roof slopes.

Adjacency relationships between roof slopes and eaves are recovered. Since ridge pixels and unclassified pixels may stand in between adjacent roof slopes, they are accounted for; the same applies to connections arising on vertices rather than along edges.

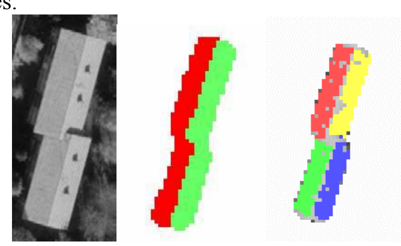

Figure 3: Segmentation of roof slopes with gradient orientation, RANSAC

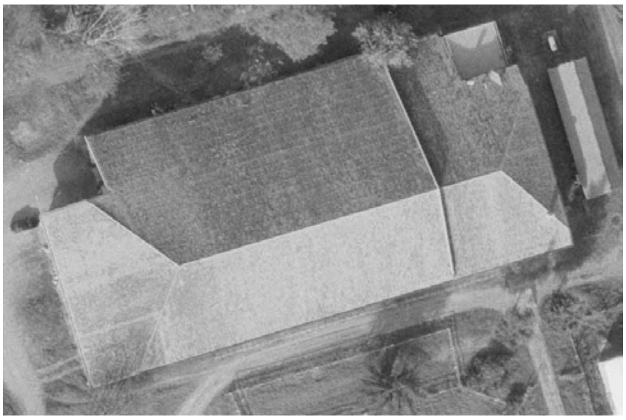

(a)

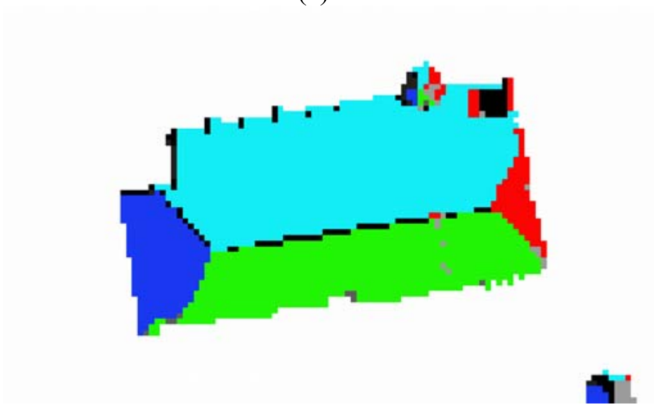

(b)

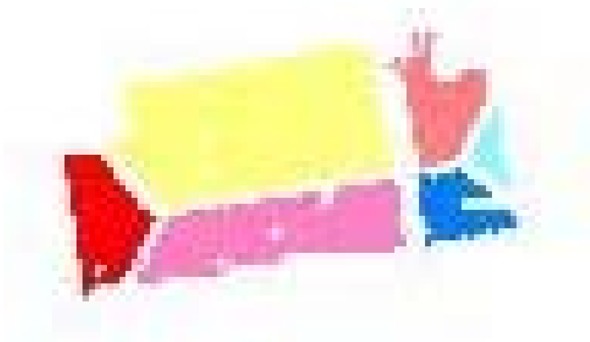

(c)

Figure 4. Segmentation based on gradient orientation (b) followed by segmentation on heights with RANSAC (c) of the building rapresented in (a)

4. CONSTRUCTION OF 3D VECTOR MODEL

The implicit building model we assume in the reconstruction procedure is a polyhedral, made of a collection of sloping planes (the roof slopes) and vertical planes (the walls), joining along the eaves lines and the roof edges. Once all planes are identified and their relationships established, intersecting terns of adjacent planes will provide with the vertices (nodes) of the roof. Linking these vertices based on a set of rules will yield the wireframe model of the roof, easily complemented by the walls, down to the nearby terrain.

Looking for roof nodes, it is worth to distinguish them in *internal* and *external*. *External nodes* belong to the roof outline, *internal nodes* are defined by three or more intersecting slopes.

Figure 5 shows the results of roof segmentation on a large sub-urban area, with the roof slopes colour-coded: it is apparent that, even in small detached houses with a noisy contour, the slopes are correctly singled out.

First *internal nodes* are determined. All terns of roof slopes labeled as adjacent are intersected. The correctness of the intersection point is tested checking if it falls into a box-shaped control volume enclosing the whole roof (Figure 8). This checks the consistency of the segmentation and helps to make up for possible missing elements, such as small slopes or step edges. If the intersection point is accepted, the proximity relationship holds and an *internal node* is established. When more than three slopes join on a single vertex, because of the inconsistencies of the solutions, distinct intersection points will be obtained. To spot such cases, internal nodes are searched for clusters, if the cluster members are close enough, they are fused and all terns are linked to the vertex. A cross check of the roof slopes linked to every pair of internal nodes yields the ridge lines.

External nodes are computed, whenever possible, by the intersection of the horizontal plane through the roof outline and pairs of adjacent roof slopes; otherwise, as the intersection points between a roof slope and two adjacent boundary walls.

The process begins with the 2D segmentation of the contour pixels in straight lines by a RANSAC strategy similar to that implemented for the roof slopes. The adjacency relationships between segments and roof slopes are reconstructed. At this point, a statistical analysis of the heights of all segments is performed to find out whether the roof border can be assumed to be horizontal.

If this is the case, each roof slope is intersected with the mean horizontal plane. Missing small roof slopes may result in invalid intersections; checks are therefore applied to insert short edges based on geometric constraints.

If not all contour pixels are horizontal, as for instance with gable roofs or with large dormers, all segments are ordered clockwise, by linking their vertices according to proximity; possible missing segments are added, to close the contour. Finally, the external nodes are computed by intersecting the vertical planes corresponding to two consecutive segments and one of the adjacent roof slopes.

First *internal nodes* are determined. All terns of roof slopes labeled as adjacent are intersected. The correctness of the intersection point is tested checking if it falls into a box-shaped control volume enclosing the whole roof. This checks the consistency of the segmentation and helps to make up for possible missing elements, such as small slopes or step edges. If the intersection point is accepted, the proximity relationship holds and an *internal node* is established. When more than three slopes join on a single vertex, because of the inconsistencies of the solutions, distinct intersection points will be obtained. To spot such cases, internal nodes are searched for clusters, if the cluster members are close enough, they are fused and all terns are linked to the vertex. A cross check of the roof slopes linked to every pair of internal nodes yields the ridge lines.

External nodes are computed, whenever possible, by the intersection of the horizontal plane through the roof outline and pairs of adjacent roof slopes; otherwise, as the intersection points between a roof slope and two adjacent boundary walls.

The process begins with the 2D segmentation of the contour pixels in straight lines by a RANSAC strategy similar to that implemented for the roof slopes. The adjacency relationships between segments and roof slopes are reconstructed. At this point, a statistical analysis of the heights of all segments is performed to find out whether the roof border can be assumed to be horizontal.

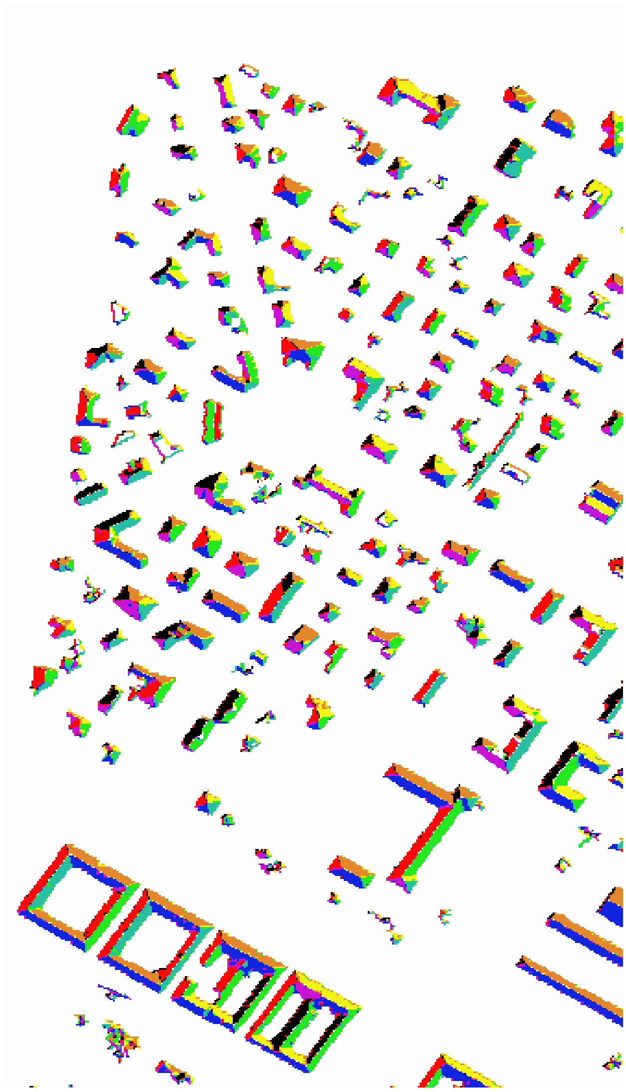

Figure 5 - Results of roof segmentation in a large sub-urban area

If this is the case, each roof slope is intersected with the mean horizontal plane. Missing small roof slopes may result in invalid intersections; checks are therefore applied to insert short edges based on geometric constraints.

If not all contour pixels are horizontal, as for instance with gable roofs or with large dormers, all segments are ordered clockwise, by linking their vertices according to proximity; possible missing segments are added, to close the contour. Finally, the external nodes are computed by intersecting the vertical planes corresponding to two consecutive segments and one of the adjacent roof slopes.

The accuracy of this piecewise approximation of the roof outlines is satisfactory only with regular shapes and long sides; it gets rather poor, because of the relatively large sampling distance, with short sides and complex shapes, sometimes leading to failure of the reconstruction. Using a horizontal plane effectively overcome these weaknesses, because roof slopes are in general more accurate than segments computed by RANSAC. Moreover, relations between sides and slopes are derived in a straight-forward manner.

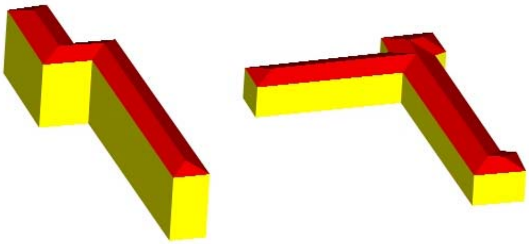

Figure 6 – The 3D models of two buildings reconstructed with the horizontal plane criterium

Once all nodes and their connections to roof slopes have been found, the *vectorization* of all the geometric elements is completed classifying them in three categories: *nodes, sides* and *polygons*. Links between these entities are stored in a *topologically encoded map*, which is the starting point to build the 3D model to be exported in different graphic formats. As a first approach, we have imported the models as 3D faces in an Autodesk environment for visualization. Each polygon of the roof is represented by a 3D surface, while external walls have been generated from eaves. In Figure 6 is presented the model of a two houses.

5. 3D VISUALIZATION

The analysis of segmentation results and of the reconstructed 3D world requires the use of a visualization tool which is able to perform efficiently the interactive exploration of data.
During the analysis of preliminary results we performed a porting to an Autodesk environment, which, in addition to the lack of an efficient terrain representation technique, soon pointed out its inefficiency in a real time analysis (exploration) of large areas with a great number of buildings.
Even if 3D visualization techniques are beyond the scope of this paper, we needed an efficient tool for real-time data exploration and verification. We chose the VTP (Virtual Terrain Project) environment. VTP writes and supports a set of software tools (VTP Toolbox) and an interactive runtime environment (VTP Enviro). The tools and their source code are freely shared to help accelerate the adoption and development of the necessary technologies". In particular, VTP adopts LODs (Level Of Details) and ROAM (Real-time Optimally Adapting Meshes) techniques, which are among the most promising approaches for general-purpose full interaction with large datasets at high frame rates, and for terrain visualization, in particular.
After installing the software, we made a porting of our data to the proprietary format adopted in VTP for representing terrain and building data, the BT and the VTST file extension, respectively. Really, the VTP environment can manage (view and process) many other kinds of geospatial data, such as road and vegetation data, but at present we have not ready the relative porting.

5 OUTLOOK ON THE FUTURE AND CONCLUSIONS

The 3D reconstruction of urban and sub-urban large environments from Lidar data provide accurate descriptions of the roofs of buildings but also useful information about the environment surrounding the buildings. All these processes are performed in an automated way even if the presence of thresholds sometime led to incorrect classification of some regions.
To increase photorealism, it is very important the introduction of textures. Information on the building façade can be supplied by terrestrial images; mobile mapping system may be the most efficient way to do it.
Sensor data fusion may be useful to resolve gaps in the laser data but consideration about the practicability of such way in term of costs and productivity has to be made. For instance, aerial photographs may help to improve data classification using texture as additional cue. High resolution satellite images may be used instead. The integration of laser data with numeric cartography allow to derive the building contours.
In conclusion, fusing data from multi sources is becoming more and more feasible (technology and costs are changing exponentially, but in opposite directions), but the extraction of maximum useful information from single data source still plays an important role, especially when it can be performed in an automated way.

References from journal:

Burns, J.B., Hanson A.R., Riseman, E.M. 1986. Extracting straight lines. In: *IEEE Transactions on Pattern Analysis and Machine Intelligence*, Vol. PAMI-8 n.4: 425-455.

References from books:

Nardinocchi C., Forlani G. (2001). *Building detection and roof extraction in laser scanning data*. In: Gruen, Baltsavias, Van Gool. (ed's) "Automatic extraction of man-made objects from aerial and space images", Ascona, 10-15 giugno 2001, Balkema, Lisse, The Netherlands, 2001, ISBN 90 5909, 319-330.

References from Other Literature:

C. Brenner, "Towards Fully Automatic Generation of City Models", *Int'l Arch. of Photogrammetry & Remote Sensing*. vol. XXXIII, Part B3, Amsterdam, 2000, pp. 85-92.

M. Duchaineau et al., "ROAMing terrain: Real-time Optimally Adapting Meshes", *Proc. Visualization '97*, IEEE, 1997, pp. 81-88.

M.A. Fischler, R.C. Bolles, "Random sample consensus: A paradigm for model fitting with applications to image analysis and automated cartography", *Comm. ACM*, 24(6), 1981, 381–395.

Forlani G., C. Nardinocchi, 2001. *Individuazione degli edifici e segmentazione dei tetti in dati LIDAR*. . Supplemento Speciale al Bollettino della Società Italiana di Topografia e Fotogrammetria, N. 2/2001, pp. 179-192.

N. Haala, C. Brenner, "Interpretation of Urban Surface Models Using 2D Building Information", *Automatic Extraction of Man-Made Objects from Aerial and Space Images (II)*, A. Gruen, E. Baltsavias and O. Henricsson, ed., Birkhauser, Basel, 1997, pp. 213–222.

K. Kraus, N. Pfeifer, 1998, "Determination of Terrain Models in Wooded Areas with Airborne Laser Scanner Data", *ISPRS*

Journal of Photogrammetry & Remote Sensing, vol. 53, pp. 193-203.

P. Lindstrom, V. Pascucci, "Visualization of Large Terrains Made Easy", *Proc. Visualization '01*, IEEE, 2001, pp. 363-370.

Nardinocchi C., Forlani G., 2001. *Detection and Segmentation of Building Roofs from Lidar Data*. ISPRS Workshop on 3D Digital Imaging and Modelling Applications of: Heritage, Industries, Medicine & Commercial Land, Padova, 3-4 April.

G. Vosselman, H.-G. Maas, "Airborne Laserscanning and Interferometric SAR for Detailed Digital Elevation Models", *Adjustment and Filtering of Raw Laser Altimetry Data*, Torlegard, Nelson, ed., Official Publication OEEPE no. 40, 2001, pp. 62-72.

U. Weidner, W. Foerstner, "Towards Automatic Building Extraction from High Resolution Digital Elevation Models", *ISPRS Journal of Photogrammetry & RS*, vol. 50 no. 4, 1995, pp. 38-49.

U. Weidner, "Digital Surface Models for Building Extraction", *Automatic Extraction of Man-Made Objects from Aerial and Space Images (II)*, A. Gruen, E. Baltsavias and O. Henricsson, ed., Birkhauser, Basel, 1997, pp.193-202.

Acknowledgements

This work has been performed in the framework of the project MURST COFIN_98: "Riprese con laser a scansione, integrate da GPS, per la produzione di modelli numerici finalizzati alla realizzazione di cartografia 3D e ortofoto digitali."