Pattern Analysis and Applications

10.1007/s10044-005-0018-2

Theoretical Advances

Complete classification of raw LIDAR data and 3D reconstruction of buildings

Gianfranco Forlani[1], Carla Nardinocchi[2], Marco Scaioni[3] and Primo Zingaretti[4]

(1) Dipartimento Ingegneria Civile, Università di Parma, Parma, Italy
(2) D.I.T.S., Università di Roma "La Sapienza", Roma, Italy
(3) D.I.I.A.R., Politecnico di Milano, Milano, Italy
(4) D.I.I.G.A., Università Politecnica delle Marche, Ancona, Italy

Gianfranco Forlani (Corresponding author)
Email: gianfranco.forlani@ipruniv.cce.unipr.it

Carla Nardinocchi
Email: carla.nardinocchi@uniroma1.it

Marco Scaioni
Email: marco.scaioni@polimi.it

Primo Zingaretti
Email: zinga@diiga.univpm.it

Received: 23 August 2005 **Accepted:** 12 October 2005 **Published online:** 10 January 2006

Abstract LIDAR (LIght Detection And Ranging) data are a primary data source for digital terrain model (DTM) generation and 3D city models. This paper presents a three-stage framework for a robust automatic classification of raw LIDAR data as buildings, ground and vegetation, followed by a reconstruction of 3D models of the buildings. In the first stage the raw data are filtered and interpolated over a grid. In the second stage, first a double raw data segmentation is performed and then geometric and topological relationships among regions resulting from segmentation are computed and stored in a knowledge base. In the third stage, a rule-based scheme is applied for the classification of the regions. Finally, polyhedral building models are reconstructed by analysing the topology of building outlines, building roof slopes and eaves lines. Results obtained on data sets with different ground point density, gathered over the town of Pavia (Italy) with Toposys and Optech airborne laser scanning systems, are shown to illustrate the effectiveness of the proposed approach.

Keywords Range images (LIDAR) - DTM - Segmentation - Classification - 3D city models - Building extraction

1 Introduction

Demands for digital terrain models (DTMs) [1] and 3D representations of the urban environment (3D city models) [2] are growing very much. High accuracy, high-resolution DTMs have become important in a wide range of applications in civil engineering, environment protection and planning, driven by raising standards (e.g. requirement of strict cost estimation of earthworks in road building) and by the development of improved prediction techniques for hydro-geological risk assessment and delimitation of areas subject to flooding. Besides, several tasks (e.g. urban planning and set-up of telecommunications networks) need 3D city models to be input in spatial databases, visualisation or virtual reality tools, wave propagation simulators, etc. Depending on the purpose these models are built for, several levels of detail and information can be incorporated. For example, buildings can be represented just as simple boxes with no pictorial information on their facades. In other cases, roof modelling and textures (artificial or photo textures from rectified images) may be required at least in some areas (e.g. 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 mapping textures on the DTM (colour raster maps or orthoimages are often used) to improve the visual representation. However, roof modelling, adding textures to building facades and including vegetation dramatically increase production costs, calling for as much automation as possible.

To fulfil demands of the above two products, LIDAR (LIght Detection And Ranging) [3] is probably the most effective technology, thanks to the high level of automation of its data acquisition workflow (raw data are routinely available just a few hours after landing) and to the accuracy and density of the 3D data delivered. For a true breakthrough in terms of overall production times to be achieved, the automation level of LIDAR data processing must be improved: the real goal today is to develop effective filtering and classification techniques for raw data, to avoid time-consuming interactive editing.

In this paper, we present a framework for LIDAR data classification and building reconstruction from LIDAR data only. Our aims are the separation of raw data in three main classes (vegetation, bare terrain and buildings) and the geometric reconstruction of polyhedral models of buildings, achieved by identifying roof slopes and eaves lines.

The paper is organised as follows. In Sect. 2
we introduce the problems of DTM generation and building reconstruction as faced by the LIDAR technology. Section 3
describes the proposed framework for LIDAR data classification. Roof segmentation and modelling are examined in Sect. 4
together with the construction of a 3D vector model for all buildings. Experimental results and some discussions are reported in Sect. 5, before conclusions.

---

2 DTM generation and building reconstruction from LIDAR data

In this section, after a brief description of the LIDAR technology, we describe the state of the art in LIDAR data processing, an overview of our approach and its main characteristics, also in comparison to other technologies.

2.1 LIDAR technology

Airborne laser scanners determine the ground position of points with an accuracy of 10–40 cm. The main components of a LIDAR system are a scanning laser telemeter, which measures the sensor-to-ground distance, and a global positioning system (GPS) receiver integrated with an inertial measurement unit (IMU). A Kalman filtering [4] determines the sensor position and attitude at the time each pulse is emitted, so that the 3D coordinates of each echo received back from a single pulse are computed. The system may be operated from helicopters or aircrafts, depending on system characteristics and mission parameters. Combination of the aircraft movement and the scanning across flight direction by an oscillating mirror results in a swath of terrain being sampled. The point density is in the order of 0.3–1 pt/m^2 for high-accuracy DTMs and may be as high as 3–5 pt/m^2 for production of 3D city models.

Most systems record the intensity of the reflected signal, which is to some extent a measure of the surface reflectance. Attempts of using this information in data classification have been made with mixed or poor results. Since visual representation of the terrain, mainly through orthophoto production or by including 3D photorealistic models, is still a key by-product of DTM production, most LIDAR systems are now coupled with an image recording system, such as an aerial photogrammetric camera or a high-resolution digital camera. Some systems can also discriminate several returns from a single emitted pulse (multiple echoes arise, for instance, when the pulse hits a power line, then a tree and finally the terrain, each surface reflecting back a fraction of the signal). Most systems retain at least the first and the last echo: since vegetation is penetrated by a certain percentage of laser pulses, this represents an effective way either to measure the canopy volume or to classify points. This is a clear advantage of LIDAR against image-based methods in DTM generation. LIDAR looks therefore better suited and more versatile than photogrammetry for DTM production; terrain breaklines extraction, on the other hand, is still critical because the anisotropic sampling pattern results in a poor definition of planimetric features at least in one direction (normally across the flight direction). Besides, data pre-processing, strip adjustment and strip georeferencing still deserve attention to remove systematic errors [5–7].

2.2 Related work

Automatic approaches to DTM generation and building reconstruction based on radiometric images have long since been a research topic [8]. Even using multiple overlapping images and sophisticated signal and feature-based matching techniques, they still face problems both in DTM generation (e.g. steep slopes, forests and urban areas) and in building reconstruction (e.g. detached houses close to high trees and old city centres). LIDAR technology has obvious advantages in both cases. Most segmentation problems arising from shadows and occlusions can be considerably reduced using range data. Height data provide directly the shape and geometric information and are independent of illumination and texture: with sufficient ground resolution, the quality of the segmentation is better than any other obtained by automatic techniques applied to optical images. Therefore photogrammetry, to date the incumbent technology, is likely to see its role diminished to a support for data analysis and data integration. Increasing cooperation between image and LIDAR data can also be expected in data processing and analysis; this is especially true in building extraction, with the former supplying edges and the latter providing roof identification and roof surfaces, but it will likely apply to data classification as well.

The main task in automatic DTM generation from LIDAR data is the separation of terrain points from non-terrain points; further classification of non-terrain points in buildings, vegetation and other objects (e.g. bridges) may be necessary, depending on the application. Different filtering strategies have been proposed, based on deviations from a parametric surface, slope threshold, clustering, etc. Kraus and Pfeifer [9] filtered out trees in forested areas by fitting an interpolating surface to the data and using an

iterative least-squares scheme to bring down the contribution of points above the surface, so that it gets closer and closer to the lowest data points; the method is effective, but it would not work in urban areas, and negative outliers may bias the interpolating surface. Rottensteiner and Briese [10] extended the technique to also filter out buildings and to discriminate them from vegetation. With this aim, they used that interpolating method iteratively on the levels of a data pyramid, where the DTM is improved in spatial resolution from a coarse initial grid; buildings and vegetation are classified by thresholding the height differences between the DTM and the digital surface model (DSM, i.e. the mathematical surface interpolating the raw data). Axelsson [11] proposed an iterative densification of a triangular irregular network (TIN), starting from seed points and adding new ones if they satisfy slope-dependent thresholds. Vosselman and Maas [6], after defining a slope threshold as the maximum allowed height difference between two points as a function of their spatial distance, developed a slope-based filtering using mathematical morphology. These last two methods do not discriminate buildings from vegetation, depend on the size of the largest non-terrain feature (the largest building in the area) and are susceptible to negative outliers (i.e. points under the true terrain surface). Brovelli et al. [12] filtered out non-terrain points by a complex iterative procedure, based on the analysis of gradient, surface orientation and residuals from successive spline interpolations. Within a threshold-based evaluation grid, points are first divided between ground and objects, and then the latter set is further divided into buildings and vegetation. Other approaches tried a classification of data points, but with partial or no explicit reference to terrain, vegetation and buildings (i.e. without an explicit separation of laser points): based on curvature and height difference analysis, Filin [13] developed a method for clustering data points in surface categories such as low and high vegetation or smooth and planar surfaces; based on connectivity and principal component analysis, Roggero [14] clustered points using geometric descriptors such as static moments, curvature and data anisotropy.

Buildings play a very important role especially in urban environments. Photogrammetry is technically capable of providing 3D city models virtually at any scale and accuracy, but with a main drawback in its costs, particularly as far as stereoscopic plotting, requiring trained personnel, is concerned. Several groups and projects have been dealing with automatic building extraction and modelling, focusing on specific stages of the process: finding whether or not there is a building in the data set (object detection), finding its position (object location) and building up its 3D representation (object reconstruction). In interactive systems, object detection may be performed by the operator, while object reconstruction is mainly or mostly automatic, thanks to the integration of constructive solid geometry (CSG), supported by a graphic user interface (GUI), in the estimation of the building model parameters [15, 16]. Indeed, witnessing the difficulty of the task, very few (and with low automaticity) commercial software packages exist [17].

As far as automatic building reconstruction is concerned, LIDAR data have been taken into consideration since mid-1990s. As a rule, much as in image-based techniques, building reconstruction is performed by a preliminary segmentation, followed by data classification and generation of building hypotheses; each hypothesis is later checked for further evidence, measuring its consistency with a building model and reconstructing its shape and location. Haala and Brenner [18], Haala [19], Brenner [20] used ground plans to detect buildings and form roof slope hypotheses; roof modelling is achieved by segmenting a DSM in planar surfaces using the RANSAC procedure [21] and reconstructing roof topology by a rule-based approach. Weidner [22, 23] detected buildings by analysing the blobs in a normalised DSM (i.e. a DSM obtained by subtracting the DTM from the original DSM) and separated 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 the minimum description length (MDL) principle.

2.3 Approach overview

This paper presents a three-stage framework for a complete, robust and automatic classification of LIDAR data as buildings, ground and vegetation, followed by a reconstruction of 3D models of the buildings. In the first stage, the raw data are filtered and interpolated over a grid. In the second stage, first a double grid data segmentation is performed (a region-growing technique to obtain regions with a step edge along their border and a grouping of connected sets of pixels on the basis of an 8-classes partition of the height gradient orientation) and then geometric and topological relationships among regions resulting from segmentation are computed and stored in a knowledge base. In the third stage, a rule-based scheme is applied for the classification of the regions. Finally, polyhedral building models are reconstructed by analysing the topology of building outlines, building roof slopes and eaves lines. The architecture of the system is sketched in Fig. 1, where, in particular, the double relationship between "Region classification" and "Knowledge base" modules refers to the iterative nature of the classification process.

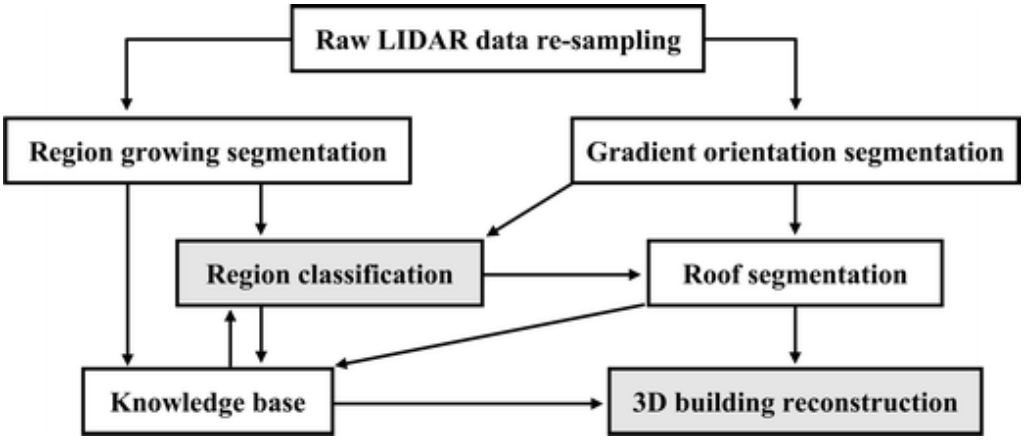

Fig. 1
Components and main relationships of the framework for LIDAR data classification and 3D building reconstruction

A first important aspect of this framework 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 recorded. Since the analysis of neighbourhood relationships is crucial to separate and label regions in building detection as well as in roof segmentation and modelling, we believe this is a key element to allow a rule-based scheme to succeed.

The rule-based scheme itself represents a second characteristic of our approach making it easily adaptable to different situations: more specific rules, depending on the environment/application, might be easily added to disambiguate the classification of some regions. In particular, this can be exploited in the case of information data fusion.

A third important characteristic of our approach is that it uses grid data to take advantage of their regularity in data processing, but it is possible to go back to raw data when necessary, for example, to extract more accurate building contours or in the estimation of roof slopes. The hierarchical structure of the framework enables reasoning about data relationships at an appropriate scale and provides the contextual information essential to increase the probability of correct classification of each data point in the final stage. For instance, when several raw data points fall into a grid cell, the interpolation algorithm that will be applied may privilege the lower points in DTM generation or the higher points in building extraction. A major motivation for this strategy is that effective filtering cannot be separated by some sort of object recognition.

Compared with other filtering techniques our strategy relies on a very simple but effective segmentation, which allows the straightforward separation of most vegetation and buildings from the terrain: in this respect, it is faster than most other approaches; it is simpler and at least as effective as Rottensteiner and Briese's technique; it is not sensitive to positive or negative outliers; as Filin's approach, though in a different way, it combines both slope and height gradient information with surface texture analysis, but with explicit classification of points; unlike slope-based filters, it is not sensitive to object size but rather, as practically all strategies, to point density; excessive fragmentation of the segmentation may result in more unclassified regions, otherwise normally less than 4%. However, a comparison of segmentation results is difficult because of the lack of quantitative metrics [*24*]. It is therefore very interesting (and worthwhile continuing) the initiative undertaken by the International Society of Photogrammetry and Remote Sensing (ISPRS) for a scientific competition on the comparison of existing automatic filters in extracting DTMs from point clouds on different landscapes and varying point densities [*25*]. The conclusions were that the filters performed well in landscapes of low complexity, while city areas and areas with sharp discontinuities in the bare earth still pose challenges. For this reason, in addition to the above-described characteristics of our approach, let us underline how the good results reported in the next sections are related to a wide and heterogeneous data set like the test site of Pavia (see Sect. 5.1), which includes urban, suburban and rural environments.

As far as building reconstruction is concerned, our approach has some similarity to Haala's, but it is based only on LIDAR data, rather than relying on 2D map data, thus overcoming problems with missing correspondences between reality and simplified or outdated maps. Indeed, in our test data, we found that the last available digital map was out of date not only in the new suburbs, but also in the old city centre. Though our building detection is rather effective, dispensing with map data makes reconstruction more difficult, since detecting edges from LIDAR data is difficult. Using parametric building models is in principle simpler and leads to consistent results; on the other hand, the variety of roof shapes found on our test site would have made it inapplicable to the Pavia case.

3 LIDAR data classification

In this section we describe in detail the procedures and the set of rules used by our data classification framework. Raw data are first interpolated to a grid (Sect. 3.1); then a region-growing algorithm, which finds connected regions bounded by a discontinuity (Sect. 3.2), and a data segmentation based on gradient orientation (Sect. 3.3) are applied; finally, the extraction of topological and geometric features of the regions (Sect. 3.4) makes it easier to apply a set of rules for data classification (Sect. 3.5).

3.1 Raw data interpolation

Four main issues arise in raw data interpolation: (a) outlier filtering; (b) the choice of the grid spacing; (c) what kind of interpolation to use when several data points fall into a grid cell; (d) how to handle data gaps:

(a) Outliers are intended here as data below the terrain surface or data above the DSM (i.e. points not falling on terrain, buildings and infrastructures, cars, vegetation, etc). They are dealt with in the region growing (see Sect. 3.2).

(b) The size of a grid cell defines the discretization of raster data classification, so it depends on the raw data point density. Ideally, in the context of building reconstruction we would use 0.5 or 1 m cell size, assuming in such cases densities larger than 1 pt/m^2, a sort of entry level for

building reconstruction. In practice, with anisotropic laser spot distributions like those of a Toposys instrument, deriving cell size just from point density (i.e. assuming a uniform distribution of laser footprints on the ground) may result in a relatively large percentage of empty cells, whose heights will be interpolated from nearby cells.

(c) If the cell is empty we interpolate with the median of the 8- or 24-connected neighbours. If the cell has just one point, we use nearest neighbour interpolation. With several points in the cell, the median is used unless there is evidence of a step edge; if this is the case, an average that weights more low points in DTM generation and more high points in building extraction is computed.

(d) We do not fill gaps larger than 3×3 cells; such gaps (possibly coming from occlusions or material with no laser response, etc.) are all conventionally classified as water.

Figure 2
shows the importance of the choice of the size of a grid cell in the transformation of raw data in grid data by interpolation. In particular, some details can totally disappear and the accuracy of object edges is reduced when using larger grid cells.

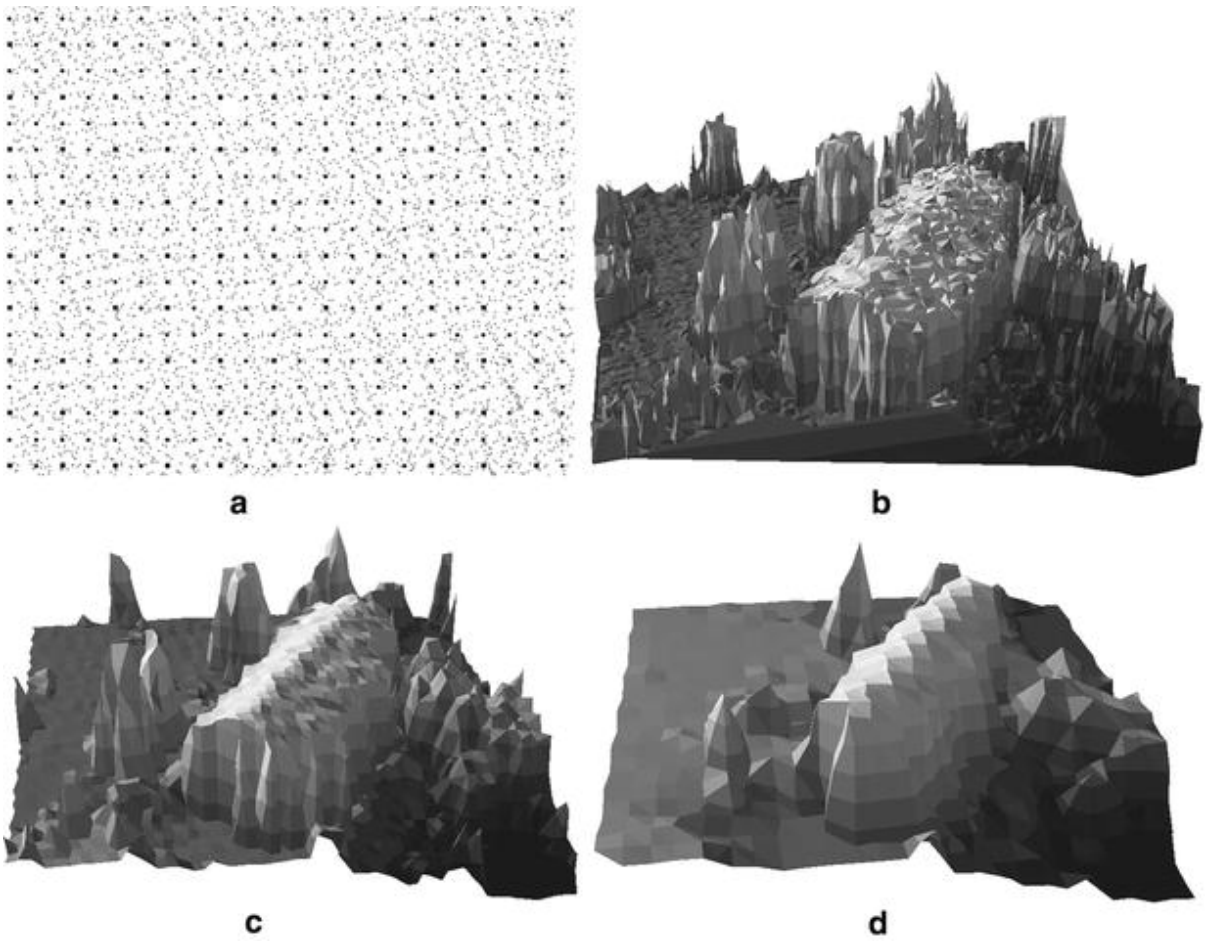

Fig. 2 From raw to grid data. **a** Raw (*grey*) points with superimposed 1 m (*black circles*) or 2 m (*black squares*) grid points; **b** triangulated DSM of raw points; DSMs after interpolation on 1 m (**c**) and 2 m (**d**) grids

3.2 Region-growing segmentation

A primary segmentation of grid data uses the following region-growing algorithm. Let begin the first region from any seed pixel; every one of its 8-connected neighbours with a height difference with respect to the seed pixel less than a threshold is joined to the region and, in turn, becomes a seed pixel for the same region. The process goes on until no more points can be joined to the region, i.e. until the region will have expanded to a continuous border with a height discontinuity larger than the threshold. The same process is started from another seed pixel to form a second region and goes on until all pixels are assigned to a region. Notice that the aggregation technique does not rule out large discontinuities within a region. The algorithm expands in all directions where the slope is smaller than the threshold; therefore, if there is a path smooth enough, it can for instance go round very steep terrain breaklines, coming to the top from behind; as it is apparent from Fig. 3, the steep dome of a church falls in the same region. Indeed, unless the DSM features closed breaklines steeper than the threshold, the procedure will include all points in a single region. With a grid size of 1 m, the threshold was set to 0.5 m after some trials; this value proved to be very effective in sorting out buildings and trees from the terrain. A higher threshold would make it possible for the algorithm to climb over buildings from nearby trees and also to include in the same region several adjacent buildings, making their subsequent reconstruction more complex. A lower threshold would cause an excessive fragmentation, which, in the end, would increase the number of unclassified regions. Since in most cases the terrain will anyway end up split in several regions (bordered, for instance, by dense tree rows, buildings or rivers), the labelling of regions corresponding to terrain will be the major task of data classification. Besides, this preliminary segmentation does not separate all instances of terrain, vegetation and buildings: a region may enclose objects of different classes (for instance, a building can be enclosed in the same region with vegetation), which must then be separated. This is achieved both by refining the segmentation on the basis of the orientation of the height gradient and by applying a set of rules, based on a geometric and topological description of the segmented regions, during the classification process.

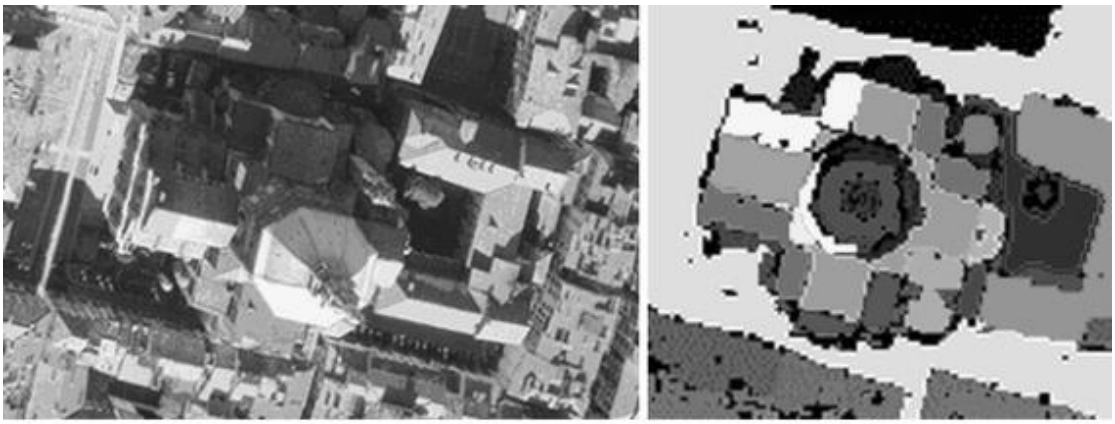

Fig. 3 Region-growing segmentation of Pavia's dome: **a** aerial image and **b** segmentation results (1 m grid)

3.3 Gradient orientation segmentation

For each pixel the height gradient orientation is computed using the Roberts operator. Connected sets of pixels belonging to the same partition of the orientation space [26] are grouped and labelled (an 8-classes partition is used). This information is then overlaid with the primary segmentation, mainly for two reasons: to discriminate between vegetation and buildings and to identify roof slopes. In fact, by analysing the distribution of gradient orientation in each region resulting from region-growing segmentation it becomes easier to distinguish vegetation from building roofs even when they have comparable heights: contrary to what happens in regions corresponding to vegetation, roof slopes,

despite noisy data, are always clearly defined and grouped in the same class. For example, this can be seen in Fig. 4
by comparing the results of gradient orientation segmentation in the upper left part, characterised by high trees, with those in main roofs and towers of Pavia's castle.

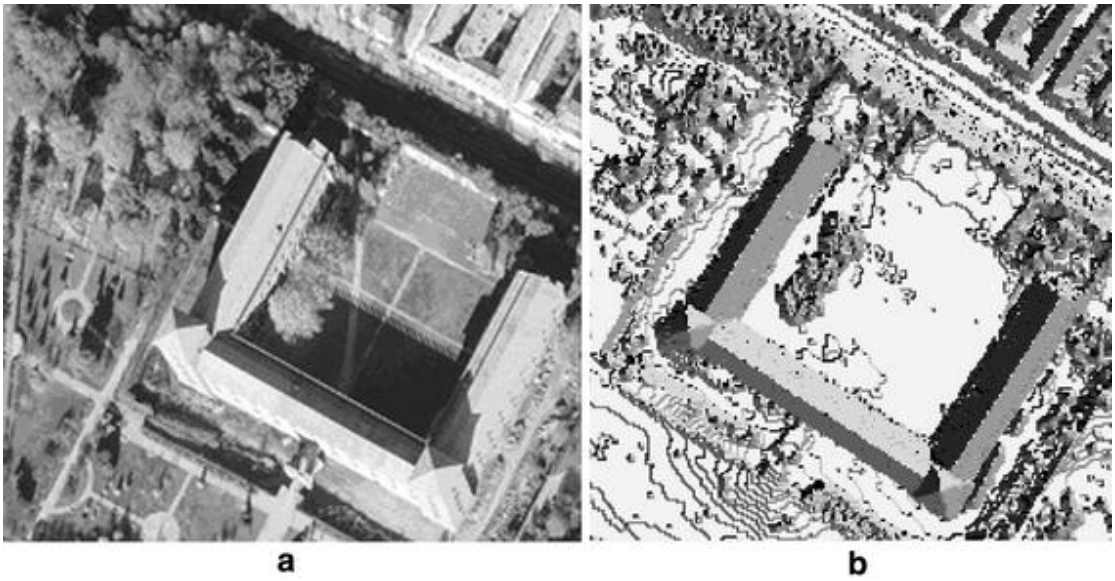

Fig. 4 Gradient orientation segmentation of Pavia's castle: **a** aerial image and **b** segmentation results (1 m grid)

The integration between the two segmentation methods is guided by the primary segmentation, which also remains the primary reference for the following processing. In particular, while regions resulting from gradient orientation segmentation can be split or merged, regions resulting from region-growing segmentation can be only split, but not merged. If, for example, a building with a gable roof is adjacent to a lower building with a one-slope roof at its left and to trees with comparable heights at its right, the results of the two segmentations could be those sketched in Fig. 5: data classification (see Sect. 3.5) based on the primary segmentation detects the two buildings, but probably includes trees together with the higher building; if the two leftmost roof slopes have slightly different orientations the gradient orientation segmentation includes them in a unique region, while vegetation is separated in many regions; at last, while the leftmost region is correctly recognised only on the basis of the primary segmentation, three sub-regions are recognised during roof slope extraction (see Sect. 4.1.1), after partitioning the region corresponding to the higher building on the basis of results of gradient orientation segmentation.

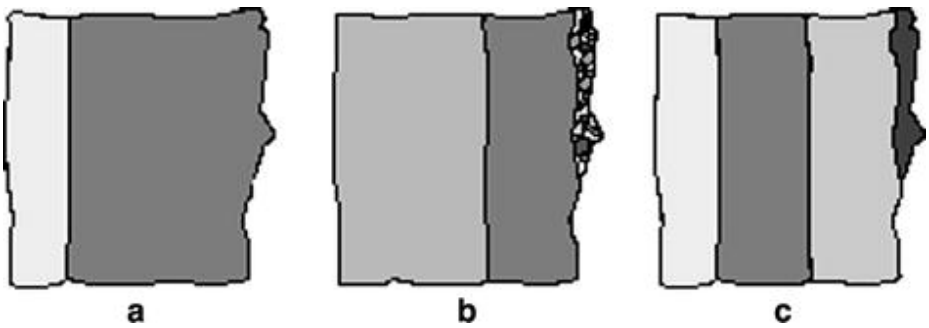

Fig. 5 Integration of segmentation results: **a** region growing; **b** gradient orientation; **c** final classification

3.4 Geometric and topological description of the regions

A geometric and topological description of the regions is a key aspect of our approach. Basic statistical parameters are computed for each region resulting from the primary segmentation. The topological information can be summarised in two graphs, the graph of external adjacencies (*EAG)* and the graph of heights *(HG*), and in a set, named the useful external border (*UEB*) of a region, defined as the set of pixels that surround the region and exist (do not belong to the image external border).

Let *A* be the set of *n* simply or multiply connected regions. An adjacency relationship (*AR)* between two regions A_i and $A_j \in A$ exists if and only if the *UEB* of A_i and A_j have a non-empty intersection. *AR* is reflexive and symmetric. If the *UEB*
is visited clockwise (so that the region is always to the right of the border), we can say that A_j has an external adjacency relationship (*EAR)* with A_i when A_j is always on the left of the external border of A_i. According to this definition, *EAR*
is transitive but neither reflexive nor symmetric. The unique difference between our *EAG* and a usual adjacency graph (in which nodes represent the regions and arcs, bi-directional or, equivalently, not oriented, the adjacency relationships) is that the arcs must be oriented: if A_i *EAR* A_j and A_j *EAR* A_i are both true (this is not mandatory) the two nodes must be connected by two opposite arcs, not by one bi-directional arc.

HG has the same number of nodes as *EAG*. The arcs of *HG* are oriented from the node corresponding to the region with the higher mean elevation along the common border towards the lower one; the difference between the mean heights is an attribute of the arc. In particular, an arc between two nodes exists only if *AR* is true.

Figure 6
reports an example of segmentation results as much realistic as representative of all typical relations among regions. Region 1 is the terrain, region 6 a courtyard, 4 is a small region and the remaining regions are buildings. The *EAG* and the *HG* of all these regions are shown in Fig. 7. This example shows clearly how the above-defined topological information can be used in data classification. According to the rules in Sect. 3.5, the algorithm is not misguided by the fact that regions 2, 4, 8 and 9 are all surrounded by both lower and higher regions and it correctly classifies them. Labelling of regions 1 and 6, on the one hand, and of regions 3, 5 and 7, on the other hand, is straightforward because their nodes in *HG* have, respectively, ingoing and outgoing arcs only.

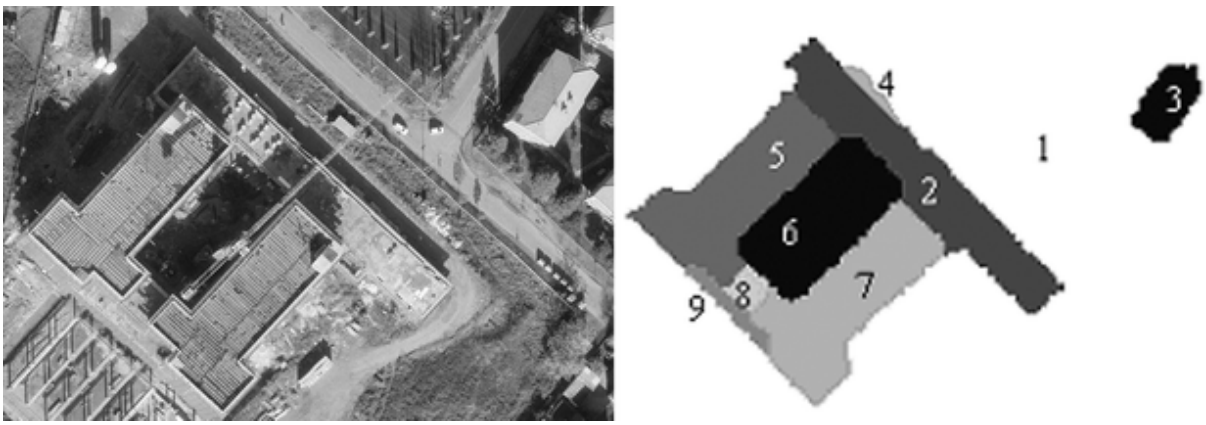

Fig. 6 A typical result of the segmentation process

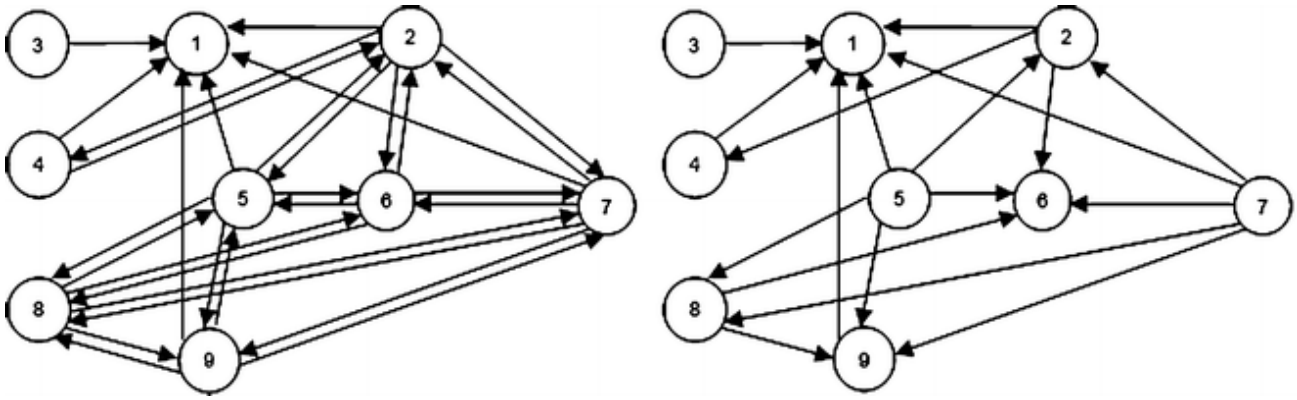

Fig. 7 The *EAG* (on the *left*) and the *HG* (on the *right*) of the regions shown in Fig. 6

3.5 Grid data classification

The hierarchical application of a set of rules, based on geometric and topologic properties of regions resulting from the segmentation, assigns the regions to one of several classes: bare terrain, buildings, vegetation, courtyards, water (or data gaps), narrow regions and noise. The process starts with the identification of vegetation and noise; then the main terrain regions and the higher buildings are extracted. At this point, some criteria (described below) to detect buildings are applied iteratively and alternatively to other rules for the detection of fragmented terrain. Anyway, this alternation between building and terrain detection is usually necessary only for complex areas, such as city centres with many adjacent buildings at different heights or residential zones where vegetation produces a high fragmentation of terrain.

3.5.1 Vegetation, narrow regions and noise

The identification of vegetation and noisy data is straightforward, owing to the characteristics of our region-growing algorithm: most points on trees fall in adjacent but distinct regions, typically 1–3 pixels large, as do proper outliers (see Sect. 3.1). Small regions are defined as those containing less than 4 pixels: if they are adjacent to other small regions, they are grouped in a single region, which is classified as vegetation; if they are isolated, they are considered as noise. When both first and last pulses are available, if pixels have differences between the height of the first and last pulses larger than 20 cm (i.e. about two times the height noise level of the sensor) they are classified as vegetation. Narrow regions are defined as those with elongated shapes; they are often found on large trees, fences, dykes, balconies, shelters, etc. Currently, we are not yet able to discriminate whether a narrow region belongs to an artificial or natural structure.

3.5.2 Terrain and courtyards

The region with the highest score in terms of geometric and shape descriptors (maximum number of pixels and maximum bounding rectangle) as well as topological relationships (maximum number of *AR* and outgoing arcs, maximum number of adjacencies with regions classified as vegetation, minimum number of lower adjacent regions) is classified as terrain. Regions exceeding a size of 200,000 m^2 are also classified as terrain, provided all their *EAR*s do not have a mean height difference larger than 5 m (in this case they are labelled as large industrial buildings). The rationale for this is that unless a small patch of data is taken around a single, flat large building, this is always the case even in dense urban

areas. Regions with only ingoing arcs in the *HG* graph (usually very small ones) are also classified as terrain or, if their adjacent regions are all buildings, as courtyards.

The terrain may get fragmented into several regions because of “fences” acting as barriers for the region-growing algorithm; this happens, for example, with data gaps, tree rows or large hedges and sometimes with very steep terrain. Looking through the *EAG* we can pick unclassified regions and find out whether the height difference across the region border is less than a threshold of 1.5 m. If this is the case, they are also considered terrain region, since the threshold is less than anything likely to be called a building, but enough to overcome rows of trees and bushes.

3.5.3 Buildings

Regions corresponding to buildings are labelled in two stages. In the first stage, rules providing a set of regions provisionally classified as buildings are applied. A first rule defines elevated regions as those with only outgoing arcs in the *HG*; they will include single buildings (i.e. a whole building enclosed in a single region) or building parts higher than those nearby (e.g. building blocks made of adjacent buildings of different height or superstructures, like turrets or cabins, in large buildings). Afterwards, to catch buildings with intermediate elevations with respect to terrain and elevated regions, the relationships in *EAG* and *HG*
between unclassified and already classified regions (terrain and elevated regions) are iteratively examined according to the following two criteria:

1. An unclassified region is labelled as a possible building if it is adjacent to at least one elevated region and is higher than all its unclassified neighbours.
2. An unclassified region is labelled as a possible building if it is adjacent to at least one terrain region and is higher than the terrain and all its unclassified neighbours by a threshold of 1.5 m.

Of course, some of the regions classified as buildings so far may represent either terrain or vegetation: in most cases classification errors occur on high trees with dense canopy. Whether these regions are really buildings is decided in the second stage, based on the segmentation of the gradient orientation. Each true building region will enclose one or more plane surfaces (regular sub-regions defined by at least four co-planar data points). Ideally, each roof slope should correspond to one regular sub-region, but due to noise and maybe overhanging trees, it may result in fragmented sub-regions and unconnected pixels (adjacent pixels belonging to different partition classes). Therefore, if the area of all regular sub-regions adds up to less than 30% of the candidate building region, this is classified as vegetation; the same applies if there is no evidence of roof slopes.

3.5.4 Unclassified regions

Some regions may remain unclassified. In most cases they are small terrain patches enclosed by dense vegetation, terrain breaklines or buildings.

---

4 Building reconstruction

All building regions are passed to the building reconstruction module. First, plane surfaces representing roof slopes and building outlines are extracted; then the topology of roof slopes and building outlines is

reconstructed. Finally, based on geometric and topological information on the roof structure, the 3D reconstruction of buildings is performed.

While good and robust identification of roof slopes is normally achieved, building outlines are harder to derive by just grouping border pixels; this results sometimes in poorly defined edges. Besides, since the building model assumes a volume bounded only by planar surfaces, buildings with domes, barrel vaults and so on might be labelled as such, but they will not be reconstructed.

4.1 Roof segmentation

Pixels within a building region are assigned to three classes: roof slopes, roof outlines, noise (small roof superstructures, such as small dormers, chimneys and lift-boxes, and vegetation). Roof slopes and points labelled as noise are obtained from the computation of gradient orientation. Contour points of a building region are derived as the union of its intersections with every *UEB* of regions with an *AR* with the building region itself.

4.1.1 Roof slope extraction

For every building region, the segmentation based on gradient orientation (see Sect. 3.3) has provided sub-regions corresponding to plane surfaces. To obtain the parameters of each plane, a least-squares fit is computed using all grid points of the sub-region. If the maximum residual is below a threshold of 15 cm, the estimated parameters of the plane are accepted; otherwise a more robust method is used to further subdivide the sub-regions. For example, two planes may have been grouped in the same sub-region in the case of parallel roof slopes or roof slopes with the same orientation but a different slope angle or just a slightly different orientation. To discriminate such cases, an iterative segmentation with the RANSAC algorithm is applied, selecting the plane with the largest consensus and discarding its points in the next iteration, until only noisy isolated points remain (Fig. 8).

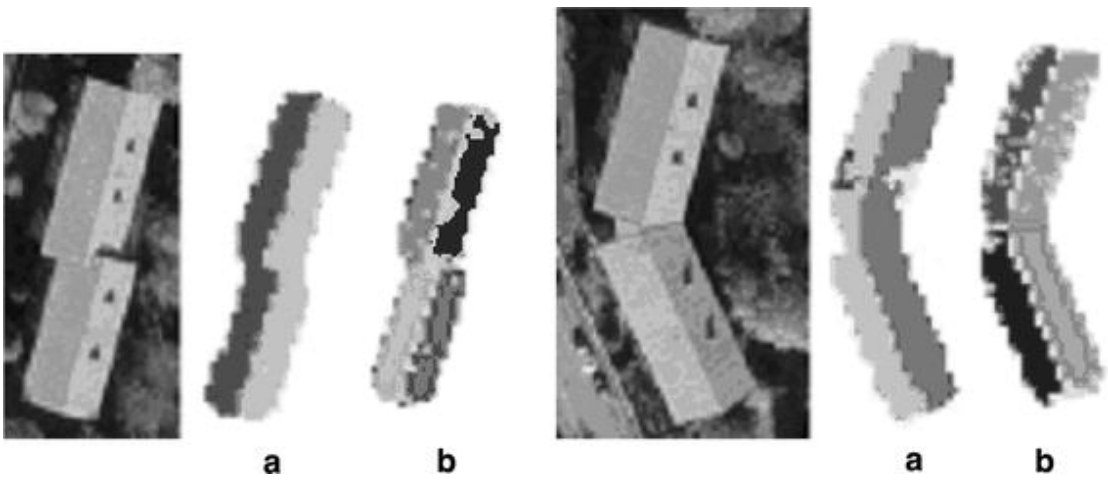

Fig. 8 Examples of erroneous roof slope segmentation by gradient orientation (**a**), corrected by RANSAC algorithm (**b**)

4.1.2 Roof outline extraction

Roof contour pixels of building regions must be segmented to provide a vector representation of eaves. To extract roof outlines we developed two methods, depending on the distribution of the heights of the

roof contour pixels. If a robust statistical analysis confirms that they belong to the same horizontal plane, each roof slope is intersected with the horizontal plane that best fits the contour pixels; in this way eaves are defined by straight lines. Otherwise, if no horizontal plane fits the data acceptably, a 2D version of the RANSAC algorithm implemented for roof slopes is applied and the contour pixels are segmented in straight line segments. The quality of this second segmentation method is normally good or acceptable with rectangular, L and U shapes. With more complex building shapes, such as in small or one-family houses with many short sides at right angles, results are less satisfactory, since details are difficult to capture with 1 m resolution. Due to the higher redundancy, the estimates of the coefficients of the planes fitting the roof slopes are more accurate than those of the straight lines of the 2D RANSAC segmentation. As it was also confirmed by applying, when suitable, both methods to the same building, the intersection with the horizontal plane leads to a more regular outline.

If any outline segment is missing, in some cases the contour can be completed by inference. For example, if the segmentation yields only three sides of a hip roof, a fourth side is added to complete the outline, provided the new side is geometrically consistent with the orientation of the nearby roof slope(s). With both methods, the endpoints of a contour segment are found by projecting its pixels on the straight line; the segments are then ordered clockwise, based on the proximity with neighbouring endpoints. A wall (a vertical plane through an outline segment) is associated with every segment of the contour.

4.1.3 Roof topology reconstruction

A topological description of the roof is built in a manner very similar to that used for regions and stored in the *Roof Adjacency Graph* (*RAG*). Similar to *EAG*, *RAG* captures the adjacency relationships between all the roof segments, i.e. roof slopes and line segments (see Figs. 10 and 11 for two examples). The associated data structure records every pixel of the *UEB* of each segment, so that the extent of the adjacency can be measured, as well as the “inner border” of the segment, i.e. the segment pixels themselves. Besides, we are examining the helpfulness of a graph of roof heights (*RHG*, in strict analogy with the *HG*) in highlighting undetected vertical planes (discontinuities) between roof slopes. The adjacency between roof slopes and contour segments is recorded only if confirmed by at least 3 pixels: this normally prevents from highlighting the “point-wise” adjacency of edges and roof slopes (such as that of “roof slope 1” with “side 4” in Fig. 9). On the other hand, the pixels of the same segment may belong to more than one roof slope, as in the gable roof of Fig. 11, so that the correspondence among roof slopes and contour segments may be not unique.

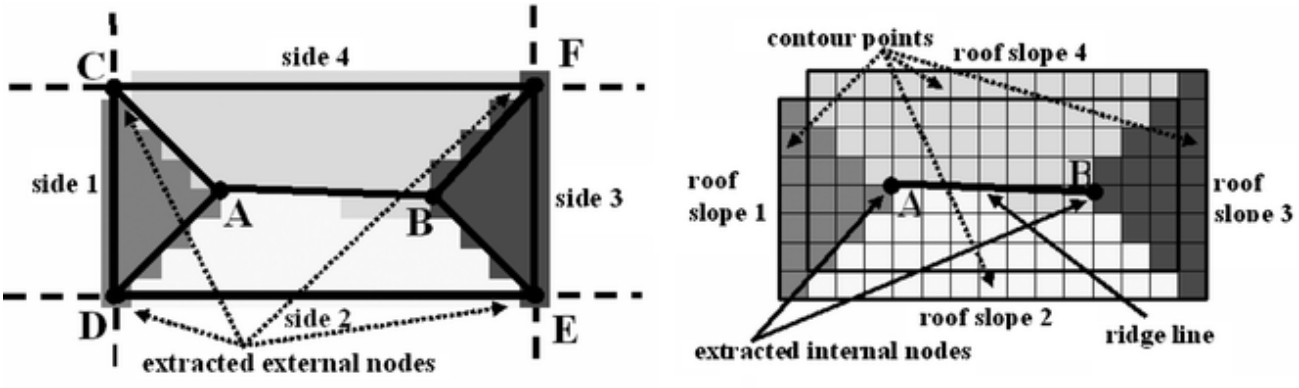

Fig. 9 Example of the extraction of *internal nodes* (on the *left*) and *external nodes* (on the *right*) of a building

4.2 Construction of the 3D vector model

The implicit building model that we assume in the reconstruction procedure is a polyhedral, made up of a collection of sloping planes (roof slopes) and vertical planes (walls), joining each other along the eaves, roof hips and ridges. The wireframe reconstruction of buildings is carried out in two steps: the identification of all vertices (*nodes*) of the roof, distinguished in *internal* and *external*, and the detection of edges connecting either roof nodes or the roof nodes to the nearby terrain. Roof nodes are computed, using the information on roof topology, as the intersection of adjacent terns of planes (either sloping or vertical): as it can be seen from Figs. 10 and 11, each triangle (tern) in the *RAG* corresponds to a uniquely determined node, while quadrilaterals (two terns with at least two common planes) correspond to nodes with multiple solutions.

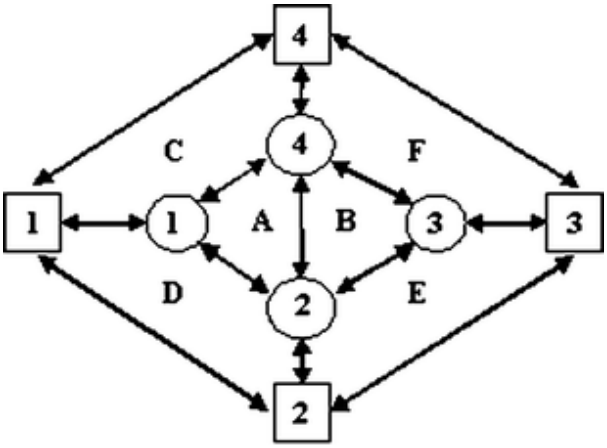

Fig. 10 *RAG* of the roof in Fig. 9. Roof slopes are in *circles*, walls in *squares*. Internal and external nodes are the *polygons*
in the graph. Edges (hips, ridges, eaves) from a node are found looking for adjacent polygons in the graph (e.g. as the polygon associated to node C is adjacent to nodes A, D and F there will be edges CA, CD and CF in the wireframe model)

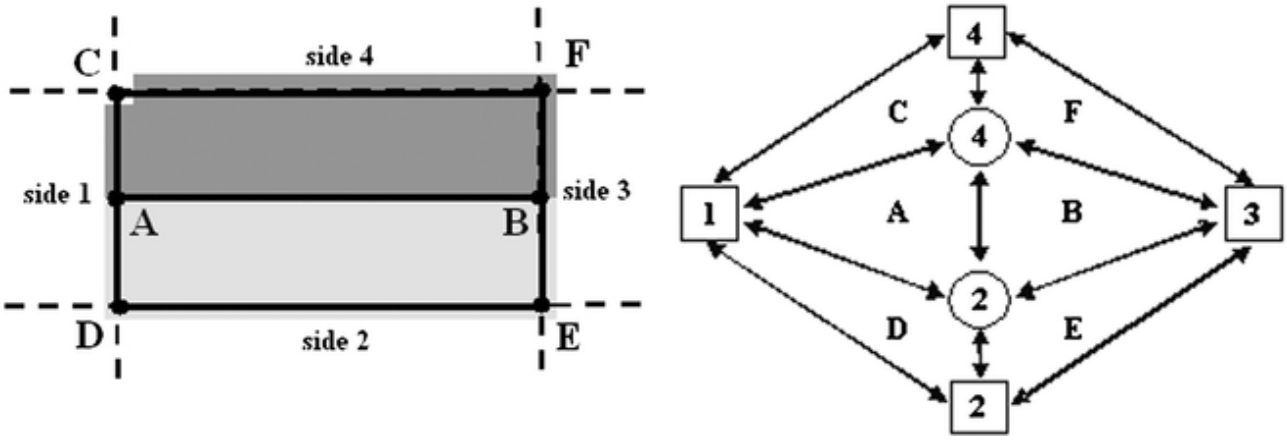

Fig. 11 A gable roof (on the *left*) and its *RAG* (on the *right*)

4.2.1 Extraction of internal nodes

Internal nodes
are defined by the intersection of terns of adjacent roof slopes. A node is computed for each feasible tern in the *RAG*; in the example of Fig. 9 there are two terns of connected roof slopes, as shown in Fig. 10: the tern (1, 2, 4) and the tern (2, 3, 4). They provide the nodes A and B. When more than three slopes join on a single vertex, because of inconsistent solutions, distinct intersection points will be obtained. To spot such cases, distances between internal nodes are computed and, if below a threshold, the nodes are fused together: all terns are therefore linked to the new vertex. To check against false internal nodes, the intersection point must fall into a *box-shaped control volume* enclosing the whole roof (Fig. 12).

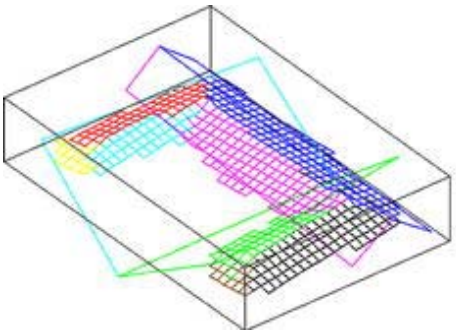

Fig. 12 The *box-shaped control volume* enclosing a roof; a valid (internal to the box, in *blue*) and an invalid (external to the box, in *green*) intersection between roof slopes are shown

4.2.2 Extraction of external nodes

External nodes
are defined either by the intersection of two adjacent contour segments (walls) with a roof slope or by the intersection of pairs of adjacent roof slopes with the contour wall (if any) common to both. Figure 9 shows an example of the former case with multiple solutions (a quadrilateral in the *RAG*): a pair of consecutive edges (e.g. 1 and 2) is intersected first with “roof slope 1” and then with “roof slope 2”. While the *x*, *y* coordinates of node D will be the same, values z_1 and z_2 will generally differ, due to the uncertainty in the estimation of planes of both roof slopes. If the difference is under a given threshold, their mean value will define the elevation of the external node D. Otherwise, this means that there is a vertical discontinuity between roof slopes 1 and 2, which are not truly adjacent; in this case the 3D model will feature a vertical face connecting both roof slopes, while a pair of vertically aligned external nodes will be introduced. In a gable roof (Fig. 11) there are no internal nodes and the polygons in the *RAG*
are only made of triangles, i.e. there are no multiple solutions. Notice that nodes A and B would not be captured by the 2D RANSAC segmentation.

4.2.3 The wireframe model

Once the coordinates of all roof nodes are available, their mutual connections (representing the roof eaves, ridges and hips) must be found. This again can be drawn from the *RAG*, since all pairs of internal and external nodes sharing two nodes of the *RAG* (either slopes or walls) are connected by an edge roof. In the example of Fig. 9, node A is connected to roof slopes 1, 2 and 4, while node D to slopes 1 and 2, resulting in a hip (AD) between roof slopes 1 and 2. By this method all the segments defining the roof are found and labelled as *external edges* (connecting a pair of external nodes), *internal ridges* (connecting a pair of internal nodes) and *hips*

connecting internal and external nodes. Finally, to complete the wireframe vector model of the whole building, vertical lines connecting each node of the external edges to the ground must also be established. While every external node has already its own height (see Sect. 4.1.2), that of the ground has to be derived from the *UEB*. Currently there is no refinement of the wireframe shape to impose constraints of symmetry, orthogonality, etc. Therefore, the quality is sometimes poor, mainly due to discretization effects, especially in case of short edges.

Though the vectorization of all geometric elements of the building is now completed, the visualisation of the reconstructed 3D model in a computer graphics environment requires the definition of the building external surfaces for rendering or for application of photorealistic materials and the like. Thanks to the topological information stored in the *RAG*, the roof polygons can be extracted: each side is defined by a pair of nodes connected to a common roof slope; a polygon is found by grouping all edges linked to a given roof slope; vertices of the polygon will have the coordinates of respective nodes. Likewise, a polygon representing a wall is derived from each external edge. We have imported all the polygons as 3D faces in an Autodesk environment to visualise the reconstructed models. Figure 13d shows the reconstruction of a large area of the Pavia town overlaid on the computed DTM.

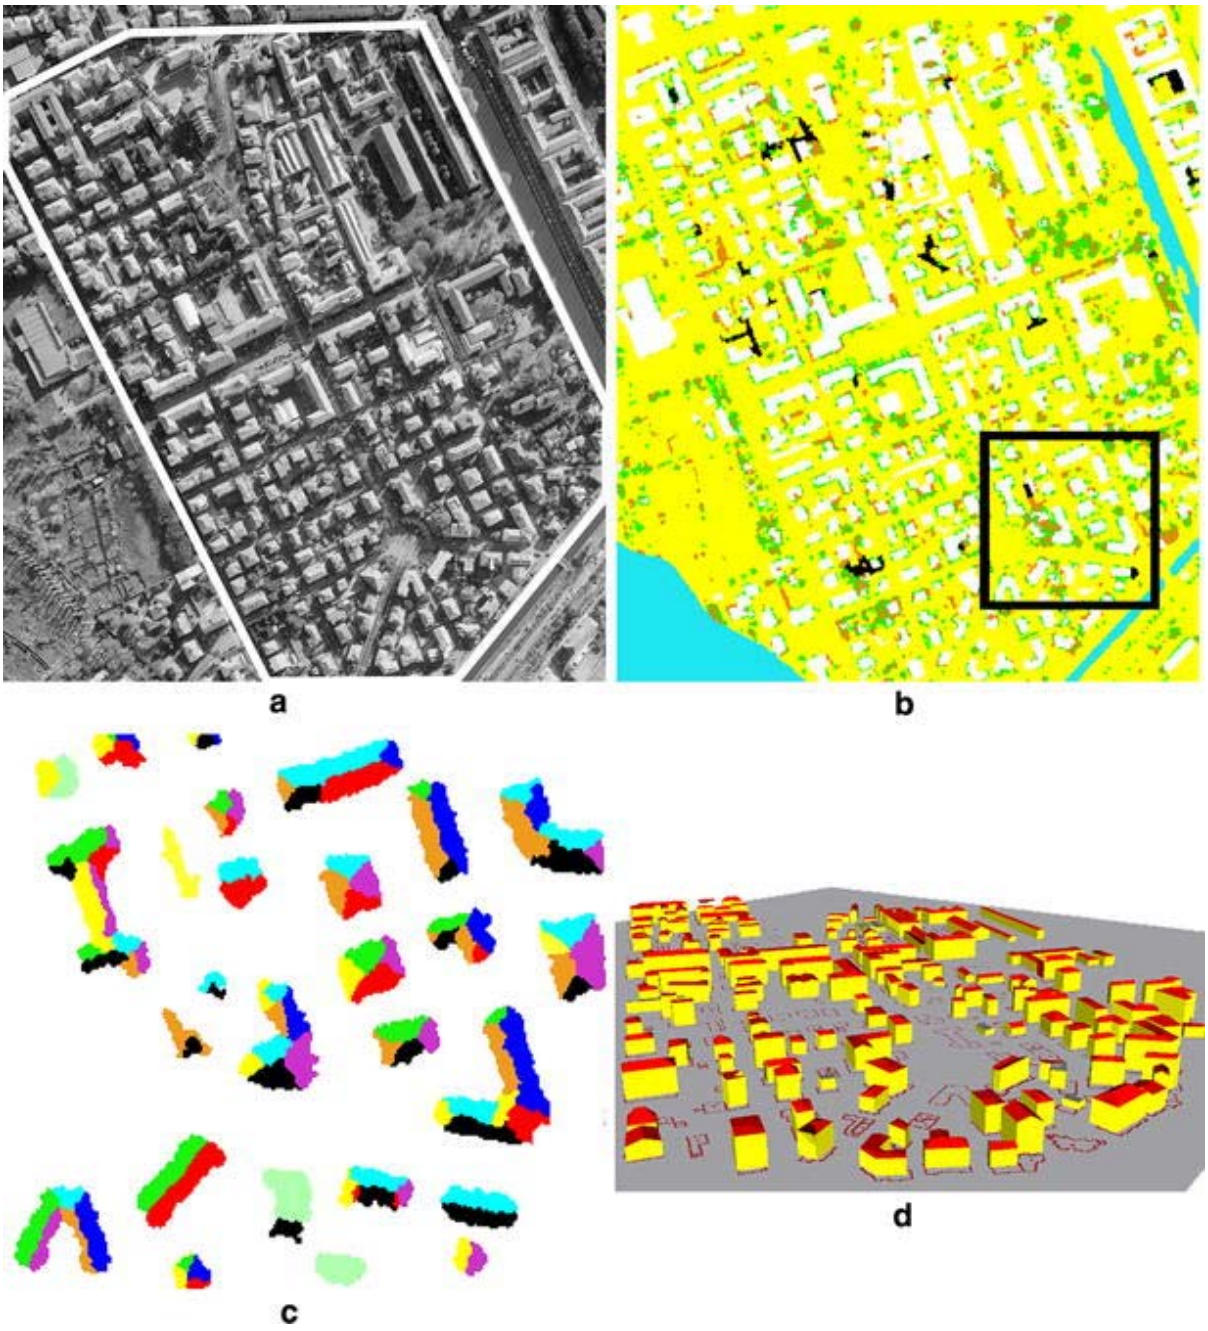

Fig. 13 Results on a residential suburban area of Pavia. **a** Aerial photo; **b** region classification: buildings (*white*); terrain (*yellow*); vegetation (*green*); narrow regions (*orange*); water (*light blue*); regions left unclassified (*black*); **c** zoom in of the square area bordered in *black* in **b**: roof slope segmentation after noise removal; *colours* show membership to the 8-classes partition of gradient orientation; **d** 3D models of buildings, superimposed to the digital map, that were successfully reconstructed in the polygonal area bordered in *white* in **a**

5 Experimental results

The proposed framework has been tested on urban and rural landscape areas with mixtures of vegetation and buildings on hillsides, data gaps, large buildings, roads with bridges and embankments, densely packed buildings with vegetation between them, vegetation on river banks and many other situations. In this section, after the description of the Pavia's data set used in the experiments to validate our strategy,

we first show the results of the classification obtained on a total area of about 5 km^2 with two different ground point densities and two different instruments; then, the behaviour of the algorithm in two large areas (one in the old city centre and the other one in a suburban zone) is shown in more detail.

5.1 Pavia's 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 (Italy), with an airborne Toposys I laser scanning system [*27*] and, about 2 weeks later, with a helicopter-based Optech 1210 system [*28*]. Stereo aerial images were also gathered during the Toposys flight. The area size is about 30 km^2 and includes the old city centre (with narrows streets and very complex roof shapes), suburban areas (with trees and both high-rise buildings and detached houses), the Ticino river in the southern part of the town and the countryside (with farms, open fields and forests). The terrain is mostly flat, with a small slope towards the Ticino river.

The Toposys instrument flew several strips in East–West direction at about 800 m above ground and two cross strips at 400 m with a scan angle of 14°. Several strips were flown with larger than normal overlaps, so overall point density ranges from about 6 to 12 pt/m^2. Either first pulse or last pulse data were available from a mission. The low scan angle has the advantage of a good penetration in narrow streets between buildings, which is typical of most old Italian cities. Due to system characteristics, patterns of laser spots on a flat terrain are regular but anisotropic, with spacing of about 1.6 m across track and 0.12–0.15 m along track.

The Optech ALTM 1210 instrument flew three missions: the first one covered the urban area at about 600 m above ground with a scan angle of 30° (point density: 1 pt/m^2); the second one over a suburban area at about 650 m with a scan angle of 40° (point density: 0.4 pt/m^2) and the last mission along a railway line at about 500 m with a scan angle of 20° (point density: 2.8 pt/m^2). All Optech missions acquired first and last pulses.

5.2 Presentation and discussion of the results

The whole Toposys data set was processed at 1 m grid resolution over about 30 km^2, while two Optech data sets were combined and processed at 2 m grid resolution (point density ranging from 0.4 to 1 pt/m^2) over about 19 km^2.

The classification process outputs three classes: terrain, buildings and vegetation; terrain, apart from bare earth, includes any surface accessible to pedestrians; vegetation is intended primarily as trees or fences; buildings are intended primarily as single family houses or larger, with plane roof surfaces. Figures for the evaluation have been computed by checking seven sites (from 0.15 to 1.5 km^2 wide) in various areas: old city centre, suburban areas and countryside for both Toposys and Optech data. The classification results have been verified with digital large-scale maps, 1:5,000 aerial images and by laying the colour-coded DSM upon the classification raster image. The overall evaluation for both data sets is summarized in the 3×4 matrix of Table 1, where the unclassified regions are also taken into account. Each row contains the ranges of the percentages, with respect to the true class, of areas that were correctly classified or wrongly assigned to other classes in the processing of the seven sites. The heterogeneity of checked data makes the extension of the obtained classification results to the whole Pavia's data set reasonable.

Table 1 Results of the classification in seven sites of the Pavia data sets (about 5 km^2, both Toposys and Optech data)

True class	Classified as (%)			
	Terrain	**Vegetation**	**Buildings**	**Unclassified**
Terrain	96–98	0	0	2–4
Vegetation	3–6	90–96	1–3	0–1
Buildings	0–1	0–1	97–100	0–2

Ranges represent the percentages, with respect to the true class, of correct or wrong classification

Some remarks can be made about the behaviour of the classification algorithm. First, due to heterogeneous data and different point density characteristics, different error types occur in different zones. On the whole, correctness of the classification in each class turned out to be more than 90%; commission errors (i.e. wrong labelling of regions) are rather small, apart from classification of low vegetation. The whole checked data set showed less than 4% of unclassified regions: most of them are indeed terrain patches in countryside and residential areas, but in the densely packed old city centre, some buildings also may remain undetected. As far as the terrain is concerned, terrain patches may remain unclassified in some complex courtyards of the old city centre, between buildings and nearby trees in suburban areas and in clearings in forest areas; this behaviour seems independent of the grid size. Evidence was never found that the terrain was mistakenly classified as vegetation or building; to some extent, this may be due to the terrain smoothness in Pavia, but was also the case with rough terrain in other data sets. As far as vegetation is concerned, keeping in mind that the objective is to avoid using vegetation points in DTM generation, trees were almost always correctly labelled: as mentioned in Sect. 3.5.3, dense canopy may be confused with buildings; fences were also recognized, if higher than the region-growing threshold and sampled with enough density to be classified as narrow regions. Low vegetation, like isolated bushes, may instead be classified as terrain or noise, depending on size, point density and grid size. With first and last pulses, which are very effective at identifying vegetation, bushes are also not erroneously labelled as terrain. As far as buildings are concerned, some one-family houses with roofs almost occluded by high trees were included in the vegetation class; in some complex blocks in the old city centre, the roof of some building parts was labelled as terrain; the same happened with buildings on very steep terrain (in our case, huts with the roof leaning against the river embankment) or with car ramps leading to a parking on a roof. Penthouses and balconies may or may not be captured, depending on their size and shape with respect to grid size: they are currently classified as narrow regions and so ignored by the reconstruction process.

The behaviour of our framework in the classification of a complex zone of Pavia's centre of about 350×350 m^2, shown in Fig. 14
by an aerial photo and the triangulated DSM of raw data points, is now illustrated. Buildings are correctly recognised iteration after iteration. This can be viewed in Fig. 15, where regions not yet classified are in black and different colours are associated to regions classified at each step. In particular, Fig. 15a shows regions classified as terrain (yellow), vegetation (green), narrow regions (orange) and not yet classified (black). Figure 15b shows in white the regions corresponding to higher buildings. The deep blue regions that appear in Fig. 15c are buildings detected in the first iteration of the first criterion described in Sect. 3.5.3, while red regions shown in Fig. 15d are the result of the application of all other iterations of the same criterion. Figure 15e reports the result (violet regions) of the application of criteria for fragmented terrain detection. Finally, Fig. 15f describes the result at the end of the alternative application of building and terrain detection criteria: here all the regions classified as buildings during the previous steps are in white, terrain in yellow, vegetation in green, narrow regions

in orange and regions left unclassified in black. Overall, buildings cover about 54% of the area, terrain 34%, vegetation 3% and narrow regions 6%; unclassified regions amount to 3% (of which about half are buildings). Vegetation misclassified as buildings amount to less than 1% of the total in the area.

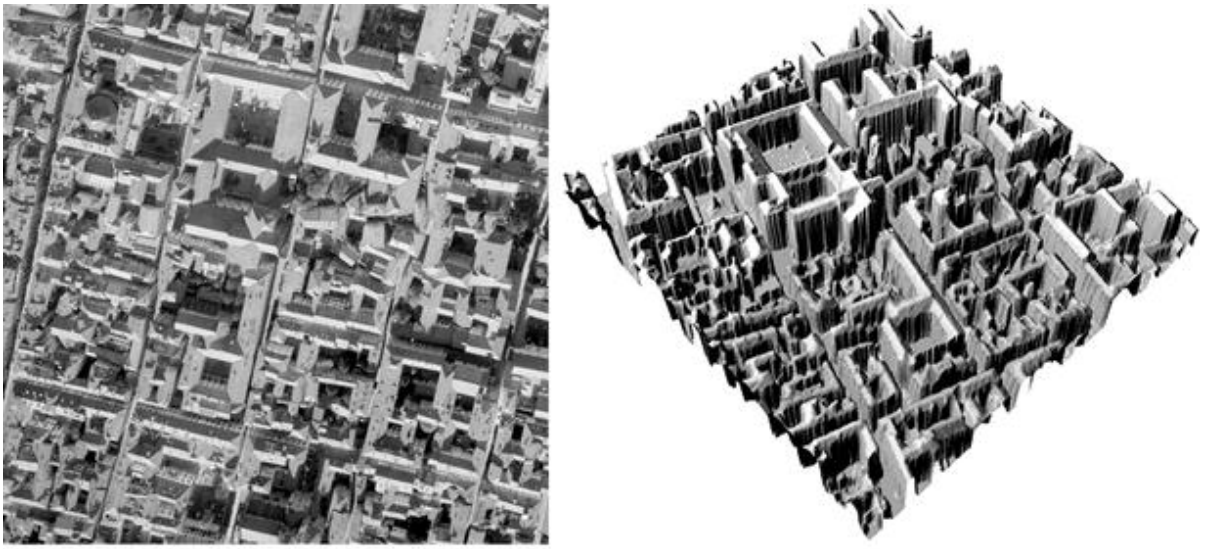

Fig. 14 Pavia's centre: **a** aerial photo and **b** triangulated DSM of the raw data points

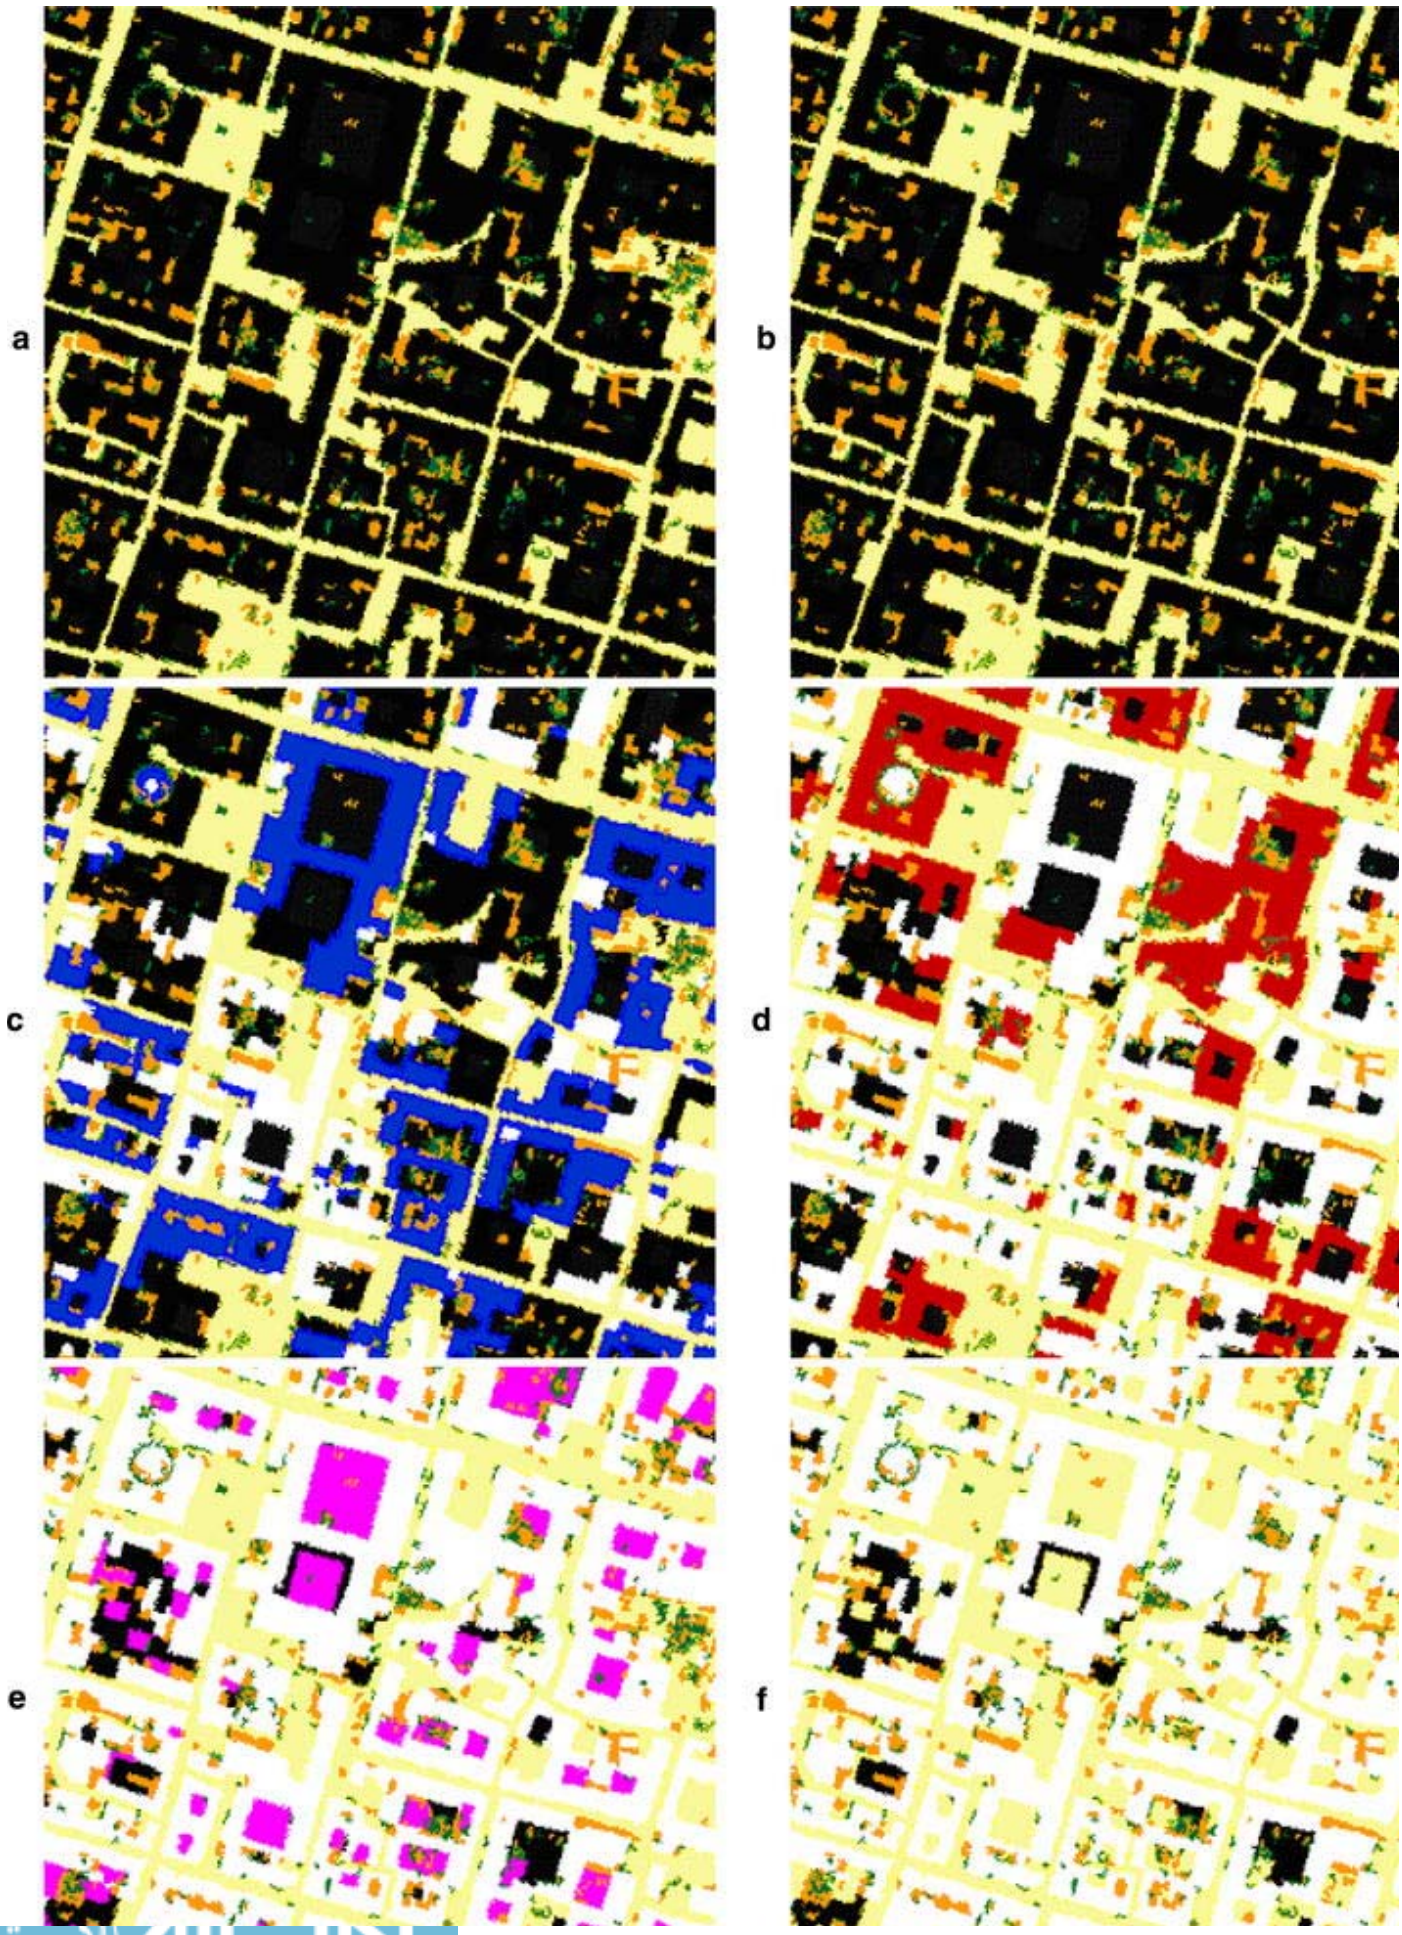

Fig. 15 Results of region classification of the zone in the centre of Pavia shown in Fig. 14: **a** initial identification of main terrain (*yellow*), vegetation (*green*), narrow regions (*orange*) and regions not yet classified (*black*); **b** detection of higher buildings (*white*); **c** detection of buildings (*deep blue*) after first application of criteria in Sect. 3.5.3; **d** buildings (*red*) detected after all other applications of criteria in Sect. 3.5.3; **e** detection of fragmented terrain (*violet*); **f** final classification: buildings (*white*), terrain (*yellow*), vegetation (*green*), narrow regions (*orange*), regions left unclassified (*black*)

Results referring to a residential, suburban area of Pavia of about 550×400 m^2, with mainly isolated buildings and trees, are shown in Fig. 13a; overall, buildings cover about 25% of the area, terrain 61%, vegetation 8% and narrow regions 5%; unclassified regions amount to 1%. The classification results (Fig. 13b) are rather satisfactory: all buildings were recognised; vegetation misclassified as building (terrain) is less than 3% (1%) of the vegetation area. There was no evidence of terrain classified either as building or vegetation.

Three-dimensional building reconstruction is less satisfactory than region classification and must be improved, though roof shapes in both the city centre and the suburbs are almost never simple gable or hip roofs as in the hypotheses. Besides, 3D building reconstruction is possible only with high-density data (more than 3 or 4 pt/m^2), which enable to work at a resolution up to 1 m×1 m. As can be seen in Fig. 13d, the 3D reconstruction fails on a relatively high number of buildings. A closer look shows that failures occurred mainly in buildings with complex shapes and many vertical discontinuities between adjacent roof slopes, which are difficult to capture (Fig. 13c). In such cases, uncorrected contour extraction and problems in segmentation and geometric modelling (i.e. computation of slope equation) of small roof slopes increase geometric inconsistencies that lead to partial reconstruction with errors. Contours might be incorrectly extracted due to the presence of short sides or of small constructions near to the main body of the building itself. In general, both problems may be attributed to the low planimetric resolution (1 m×1 m), which did not enable a correct segmentation.

Sometimes, mainly in low-rise buildings, contour definition is poor because of trees that partially cover the roof. Failures in roof slopes modelling or identification, especially when dealing with small roof slopes, lead to errors in roof topology and therefore in false or wrong computation of nodes. The criterion of the horizontal plane for contour extraction is based on the accuracy of roof plane estimation, so it fails in the presence of poorly estimated small roof slopes. Considering a numerical evaluation of failures of former and latter cases, about 70% of cases result from uncorrected contour extraction, while the remaining 30% from roof slope segmentation.

As far as execution times are concerned, the classification of the two areas (of about 350×350 and 550×400 m^2) takes less than 2 and 4 s, respectively, on an AMD Athlon 4 processor with 750 MB RAM.

6 Conclusions and perspectives

Digital terrain models and 3D city models have become more and more essential in a wide range of applications in civil engineering, environment protection and planning. A complete, robust classification of acquired data is necessary irrespective of the level of detail they are built for. In any case, due to the large amount of data to be processed, it is fundamental that the construction process is made more and more automatic.

The automaticity is the main characteristic of the framework presented in this paper, without penalising the accuracy of the built models. We combined the processes of data classification and building

extraction to avoid the need of tuning thresholds to the different environments being processed. After an initial data segmentation performed using a region-growing algorithm and complemented by a segmentation on gradient orientation, the classification, mostly as buildings, ground and vegetation, is performed iteratively by exploiting a set of rules that takes advantage of a geometric and topological description of segmented data. This strategy greatly increases the robustness of the classification, providing support for contextual reasoning. The topological structure of the segmented regions and, in particular, adjacency relationships have also been exploited in roof segmentation and roof modelling, allowing a correct reconstruction for the majority of roofs, while the rest, with relatively more complex topology, remain harder to deal with, even by increasing data resolution. Although none of the many adopted pattern analysis techniques is completely new (i.e. all of them have been separately used in other application fields), the proposed novel framework shows good results and, above all, it is easily adaptable to many engineering applications: specific rules, depending on the application, might be easily introduced, while the classification obtained in output by our framework can be used directly or fused with information coming from other sources. However, while fusing data from multiple sources is becoming more and more feasible (technology and costs are changing exponentially, but in opposite directions), in this paper we concentrated on the extraction of the maximum useful information from a single data source, because such data still play an important role, especially when they can be processed in an automated way.

Future work will be mainly oriented in the following directions. Data classification should be improved by: (1) detecting misclassified buildings included in terrain regions by a deeper gradient analysis; (2) exploiting contextual reasoning to increase the number of classes, in particular with the inclusion of bridges and overpasses, and to label some of the elongated regions; (3) comparing our manually tuned rules with rules learned by classification trees. As far as building reconstruction, developments are expected in both segmentation and modelling: (1) going back to raw data, now mainly used to get a correct classification, should be extended to improve the definition of roof edges, especially contours; (2) methods to detect discontinuities within the roof by exploiting the *RHG* (and the definition of the *RHG* itself) should be properly set up.

7 About the authors

Gianfranco Forlani, born in 1958, received the degree in Civil Engineering in 1984 and PhD in Geomatics in 1989, both from the Politecnico of Milan, Italy. In 1990, he worked under a research contract at the ETH Zürich on GPS-assisted aerial triangulation. In 1992, he became associate professor at the Politecnico of Milan. In the period 1992–1996, he was the Secretary of the ISPRS TC 1. Since 1996, he is professor of Surveying and Photogrammetry at the Department of Civil Engineering of the University of Parma. His scientific interests are in automatic and semi-automatic methods for surface reconstruction from digital images; automatic aerial

triangulation; calibration of integrated GPS/INS systems; filtering and classification of laser scanning data for DTM production and building reconstruction; mobile mapping. He is the author of 94 publications, 54 in English.

Carla Nardinocchi
received the degree in Civil Engineering in 1992 from the University of Ancona, Italy, and PhD in Geodesy and Photogrammetry in 1996 from the Politecnico of Milan, Italy. In 1998, she obtained a 2-year post-graduate research grant from the University of Parma. Since 2000, she is a full-time researcher at the Faculty of Engineering (D.I.T.S.) of the University of Rome “La Sapienza”. Her scientific interests are in automatic object recognition and feature extraction, building recognition, building reconstruction, DTM generation and terrain-point extraction from airborne laser scanner data.

Marco Scaioni, born in 1969, received the degree in Civil Engineering in 1995 and PhD in Surveying and Mapping in 1999 both from the Politecnico of Milan, Italy. In the period 2000–2002, he was assistant researcher at the University of Parma. Since 2003, he is a staff researcher at the Surveying Department (D.I.I.A.R.) of the Politecnico of Milan, where he currently holds the course of Photogrammetry and other lectures about Surveying and Mapping. His research interests first concerned the problem of automation in digital photogrammetry, then were prevalently focused on automatic building reconstruction from LIDAR data, and now concern terrestrial laser scanning, focusing in particular on problems of 3D-view georeferencing and on applications to geological field and to the monitoring of large structures deformations.

Primo Zingaretti
received the degree in Electronic Engineering in 1984 from the University of Ancona, Italy. From 1985 to 1987, he worked under annual contract at the Computer Science Institute of the University of Ancona on the implementation of computer vision and artificial intelligence software applications. From 1987 to 1990, he obtained annual grants at INRCA of Ancona for developing medical image processing systems. In 1990, he joined the University of Ancona, now Università Politecnica delle Marche, where he is currently an assistant professor at the Faculty of Engineering (D.I.I.G.A.). He is the author of over 75 papers in the fields of artificial intelligence, image processing, image understanding, perception and learning. General Chair of ECMR'05, member of IEEE and of various scientific committees, his current research focuses on applying artificial intelligence and computer vision techniques to mobile robots.

Acknowledgments This work has been partly financed under the Italian national research project COFIN no. 9808229941_002.

References

1. Weibel R, Heller M (1991) Digital terrain modelling. In: Maguire DJ et al (eds) Geographical information systems—principles and applications. Longman, London, pp 269–297

2. Shiode N (2001) 3D urban models: recent developments in the digital modelling of urban environment in three-dimensions. GeoJournal 52(3):263–269

3. Baltsavias EP (1999) Airborne laser scanning: existing systems and firms and other resources. ISPRS J Photogramm Remote Sens 54:164–198

4. Ledermann W (1980) Kalman filtering. Handbook of applicable mathematics, statistics, vol IV. Emlyn Lloyd, New York

5. Kilian J, Haala N, Englich M (1996) Capture and evaluation of airborne laser scanner data. Int Arch Photogramm Remote Sens Spat Inf Sci 31(3):383–388

6. Vosselman G, Maas H-G (2001) Adjustment and filtering of raw laser altimetry data. OEEPE Official Publ 40:62–72

7. Burman H (2002) Laser strip adjustment for data calibration and verification. Int Arch Photogramm Remote Sens Spat Inf Sci 34(3A):67–72

8. Noronha S, Nevatia R (2001) Detection and modeling of buildings from multiple aerial images. IEEE Trans Pattern Anal Mach Intell 23(5):501–518

9. Kraus K, Pfeifer N (1998) Determination of terrain models in wooded areas with airborne laser scanner data. ISPRS J Photogramm Remote Sens 53(4):193–203

10. Rottensteiner F, Briese C (2002) A new method for building extraction in urban areas from high-resolution LIDAR data. Int Arch Photogramm Remote Sens Spat Inf Sci 34(3A):295–301

11. Axelsson P (2000) DEM generation form laser scanner data using adaptive TIN models. Int Arch Photogramm Remote Sens Spat Inf Sci 33(B4/1):110–117

12. Brovelli MA, Cannata M, Longoni UM (2002) Managing and processing LIDAR data within GRASS. http://www.ing.unitn.it/~grass/conferences/GRASS2002/proceedings/proceedings/papers_by_author.html (accessed 4 April 2004). Also in Proceedings of open source GIS—GRASS user's conference, Trento, Italy

13. Filin S (2002) Surface clustering from airborne laser scanning data. Int Arch Photogramm Remote Sens Spat Inf Sci 34(3A):119–124

14. Roggero M (2002) Object segmentation with region growing and principal component analysis. Int Arch Photogramm Remote Sens Spat Inf Sci 34(3A):289–294

15. Foerstner W (1999) 3D-city models: automatic and semiautomatic acquisition methods. In: Fritsch D, Spiller R (eds) Photogrammetric week '99. Wichmann Verlag, Heidelberg, pp 291–303

16. Guelch E, Mueller H, Läbe T (1999) Integration of automatic processes into semi-automatic building extraction. Int Arch Photogramm Remote Sens Spat Inf Sci 32(3-2W5):177–186

17. Guelch E (2000) Digital systems for automated cartographic feature extraction. Int Arch Photogramm Remote Sens Spat Inf Sci 33(B29):241–256

18. Haala N, Brenner C (1997) Interpretation of urban surface models using 2D building information. In: Gruen A, Baltsavias E, Henricsson O (eds) Automatic extraction of man-made objects from aerial and space images (II). Birkhauser Verlag, Basel, pp 213–222

19. Haala N (1999) Combining multiple data sources for urban data acquisition. In: Fritsch D, Spiller R (eds) Photogrammetric week '99. Wichmann Verlag, Heidelberg, pp 329–339

20. Brenner C (2000) Towards fully automatic generation of city models. Int Arch Photogramm Remote Sens Spat Inf Sci 33(3):85–92

21. Fischler MA, Bolles RC (1981) Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography. Commun ACM 24(6):381–395

22. Weidner U, Foerstner W (1995) Towards automatic building extraction from high resolution digital elevation models. ISPRS J Photogramm Remote Sens 50(4):38–49

23. Weidner U (1997) Digital surface models for building extraction. In: Gruen A, Baltsavias E, Henricsson O (eds) Automatic extraction of man-made objects from aerial and space images (II). Birkhauser Verlag, Basel, pp 193–202

24. Hoover A, Jean-Baptiste G, Jiang X, Flynn PJ, Bunke H, Goldgof DB, Bowyer K, Eggert DW, Fitzgibbon A, Fisher RB (1996) An experimental comparison of range image segmentation algorithms. IEEE Trans Pattern Anal Mach Intell 18(7):673–689

25. ISPRS (2003) ISPRS test on extracting DEMs from point clouds: a comparison of existing automatic filters. http://www.enterprise.lr.tudelft.nl/frs/isprs/filtertest/ (accessed 4 April 2005)

26. Nardinocchi C, Forlani G (2001) Detection and segmentation of building roofs from Lidar data. In: Proceedings of ISPRS workshop on 3D digital imaging and modelling applications of heritage, industries, medicine & commercial land, Padua, Italy

27. http://www.toposys.com/ (accessed 4 April 2005)

28. http://www.optech.ca/ (accessed 4 April 2005)