



(51) International Patent Classification:
G06T 17/05 (2011.01)

(21) International Application Number:
PCT/AU2011/000017

(22) International Filing Date:
7 January 2011 (07.01.2011)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
2010200143 14 January 2010 (14.01.2010) AU

(71) Applicant (for all designated States except US): **THE UNIVERSITY OF SYDNEY** [AU/AU]; Parramatta Road, Sydney, New South Wales 2006 (AU).

(72) Inventors; and

(75) Inventors/Applicants (for US only): **DOUILLARD, Bertrand** [FR/AU]; 90A Shepherd Street, Chippendale, New South Wales 2008 (AU). **UNDERWOOD, James** [AU/AU]; 9/111 McEvoy Street, Alexandria, New South Wales 2015 (AU). **VLASKINE, Vsevolod** [AU/AU]; 13/12 Leichhardt Street, Glebe, New South Wales 2037 (AU).

(74) Agent: **GRIFFITH HACK**; Level 3, 509 St Kilda Road, Melbourne, Victoria 3004 (AU).

(81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IS, JP, KE, KG, KM, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PE, PG, PH, PL, PT, RO, RS, RU, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

(84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Published:

— with international search report (Art. 21(3))

(54) Title: EXTRACTION PROCESSES




FIG. 1

(57) Abstract: A process and apparatus for identifying a value of a parameter measured in a cell in an extraction process as corresponding to a particular object (8, 10) or terrain feature (6), comprising: for a measurement of the parameter in the cell: dividing a range of possible parameter values into voxels such that the value of the measurement lies within a voxel; and identifying the measured value as corresponding to the particular object (8, 10) or terrain feature (6) if: the voxel within which the measured value lies coincides with a range of values corresponding to a particular object (8, 10) or terrain feature (6); or there is no voxel between the voxel within which the measured value lies and a voxel that coincides with a range of values corresponding to a particular object (8, 10) or terrain feature (6) within which below a certain number of measurements of the parameter have been made.



EXTRACTION PROCESSES

FIELD OF THE INVENTION

The present invention relates to extraction, extraction processes, extraction algorithms, and the like.

BACKGROUND

Data corresponding to the geometry of an area of terrain and any natural and/or artificial features or objects of the area may be generated. For example, a laser scanner, such as a Riegl laser scanner, may be used to scan the area of terrain and generate 3D point cloud data corresponding to the terrain and the features.

Various algorithms for processing 3D point cloud data of a terrain area are known. Such algorithms are typically used to construct 3D terrain models of the terrain area for use in, for example, path planning or analysing mining environments.

The terrain models conventionally used include the Mean Elevation Map, the Min-Max Elevation Map, the Multi-Level Elevation Map, the Volumetric Density Map, Ground Modelling via Plane Extraction, and Surface Based Segmentation.

Mean Elevation Maps are commonly classified as $2^{1/2}$ D models because the third dimension (height) is only partially modelled. In these models the terrain is represented by a grid having a number of cells. The height of the laser scanner returns falling in each grid cell is averaged to produce a single height value for each cell. An advantage of averaging the height of the laser returns is that noisy returns can be filtered out. However, this technique cannot capture overhanging structures, such as tree canopies.

Min-Max Elevation Maps are also used to capture the height of the returns in each grid cell. The difference between the maximum and the minimum height of the laser scanner returns falling in a cell are computed. A cell is declared occupied if its calculated height difference exceeds a pre-defined threshold. These height differences provide a computationally efficient approximation to the terrain gradient in a cell. Cells which contain too steep a slope or are occupied by an object will be characterized by a strong gradient and can be identified as occupied. An advantage of this technique is that approximations are not made, i.e. averaging is avoided. However, this technique is more sensitive to noise than a Mean Elevation Map.

Multi-Level Elevation Maps are an extension of elevation maps. Such algorithms are capable of capturing overhanging structures by discretising the vertical dimension. They also allow for the generation of large scale 3D maps by recursively registering local maps.

Typically however, the discrete classes chosen for the vertical dimension may not facilitate segmentation. Also, typically the ground is not used as a reference for vertical height.

Volumetric Density Maps discriminate between soft and hard obstacles. This technique breaks the terrain area into a set of voxels and counts in each voxel the number of hits and misses sensor data. A hit corresponds to a return that terminates in a given voxel. A miss corresponds to a laser beam going through a voxel. Regions containing soft obstacles, such as vegetation, correspond to a small ratio of hits over misses. Regions containing hard obstacles correspond to a large ratio of hits over misses. While this technique does allow the identification of soft obstacles (the canopy of the trees for instance), segmenting a scene based on the representation it provides would not be straightforward since parts of objects would be disregarded (windows in buildings or patches of vegetation for instance).

A Ground Modelling via Plane Extraction approach is suitable for extracting multi-resolution planar surfaces. This involves discretising the area terrain into two superimposed 2D grids of different resolutions, i.e. one grid has larger cells than the other. Each grid cell in each of the two grids is represented by a plane fitted to the corresponding laser returns via least square regression. A least square error for each plane in the each grid is computed. By comparing the different error values, several types of regions can be identified. In particular, the values are both small in sections corresponding to the ground. Also, the error value of the larger celled plane is small while error values of the smaller plane is large in areas containing a flat surface with a spike (e.g. a thin pole for instance). Also, both error values are large in areas containing an obstacle. This method is able to identify the ground while not averaging out thin vertical obstacles (unlike a Mean Elevation Map). However, it is not able to represent overhanging structures.

Surface Based Segmentation performs segmentation of 3D point clouds based on the notion of surface continuity. Surface continuity is evaluated using a mesh built from data. The mesh is generated by exploiting the physical ordering of the measurements which implies that longer edges in the mesh or more acute angles formed by two consecutive edges directly correspond to surface discontinuities. While this approach performs 3D segmentation, it does not identify the ground surface.

Thus, there is a need for an algorithm for performing segmentation of 3D point cloud data that jointly provides a representation of the ground, and representations of objects.

SUMMARY OF THE INVENTION

In a first aspect the present invention provides a process of identifying a value of a parameter measured in a cell in an extraction process for extracting an object or terrain feature as corresponding to a particular object or terrain feature, the process comprising:

for a measurement of a parameter in the cell: dividing a range of possible parameter values into a plurality of voxels such that a value of measurement of the parameter in the cell lies within one of the voxels; and identifying the measured value of the parameter as corresponding to the particular object or terrain feature if: the voxel within which the measured parameter value lies coincides with a range of values that corresponds to a particular object or terrain feature; or there is no voxel between the voxel within which the measured parameter value lies and a voxel that coincides with a range of values that corresponds to a particular object or terrain feature within which below a certain number of measurements of the parameter have been made.

The process may further comprise: for a measurement of the parameter in the cell, identifying the measured value of the parameter as possibly corresponding to the particular object or terrain feature if there does not exist two different voxels which are different from the voxel within which the measured parameter value lies such that there are there is no voxel between the voxel within which the measured parameter value lies and each of the two different voxels within which below a certain number of measurements of the parameter have been made.

The process may further comprise updating a value of a function only if the value of the measured parameter is identified as corresponding to the particular object or terrain feature.

In a further aspect the present invention provides an apparatus for generating a model of terrain in accordance with the process of any of the above aspects of the invention, the apparatus comprising scanning and measuring apparatus for measuring values of a parameter, and one or more processors arranged to perform the processing steps of any of the above aspects of the invention.

In a further aspect the present invention provides a computer program or plurality of computer programs arranged such that when executed by a computer system it/they cause the computer system to operate in accordance with the process of the above aspects of the invention.

In a further aspect the present invention provides a machine readable storage medium storing a computer program or at least one of the plurality of computer programs according to the above aspect of the invention.

BRIEF DESCRIPTION OF THE DRAWINGS

Figure 1 is a schematic illustration of an example terrain modelling scenario in which a laser scanner is used to scan a terrain area;

Figure 2 is a process flowchart showing certain steps of a terrain modelling algorithm performed by a processor;

Figure 3 is a process flowchart showing certain steps of the ground extraction process of step s2 of the algorithm;

5 Figure 4 is a schematic illustration of three cells of a grid of a Mean Elevation Map;

Figure 5 is a process flowchart showing certain steps of the object segmentation process of step s4 of the algorithm;

Figure 6 is a schematic illustration of a first cell, a second cell, and a third cell of the grid and the range of height values assigned to each of these cells in a Mean Elevation Map;

10 Figure 7 is a schematic illustration of the first cell, the second cell, and the third cell of the Mean Elevation Map grid, after performing step s28;

Figure 8 is a schematic illustration of the laser scanner scanning an object that is hidden behind a further object;

15 Figure 9 is a process flow chart showing certain steps of an iterative segmentation algorithm;

Figure 10 is a schematic block diagram showing certain details of a further implementation of the process of Figure 9;

Figure 11 is a process flow chart showing certain steps of a ground model updating process; and

20 Figure 12 is a process flow chart showing certain steps of a process for determining whether a new measurement corresponds to the ground.

DETAILED DESCRIPTION

25 The terminology "terrain" and "terrain features" are used herein to refer to a geometric configuration of an underlying supporting surface of an environment or a region of an environment. The terminology "object" is used herein to refer to any objects or structures that exist above (or below) this surface. The underlying supporting surface may, for example, include surfaces such as the underlying geological terrain in a rural setting, or the artificial support surface in an urban setting, either indoors or outdoors. The geometric configuration of other objects or structures above this surface, may, for example, include
30 naturally occurring objects such as trees or people, or artificial objects such as buildings or cars.

Some examples of terrain and objects are as follows: rural terrain having hills, cliffs, and plains, together with object such as rivers, trees, fences, buildings, and dams; outdoor urban terrain having roads and footpaths, together with buildings, lampposts, traffic lights, cars, and people; outdoor urban terrain such as a construction site having partially laid foundations, together with objects such as partially constructed buildings, people and construction equipment; and indoor terrain having a floor, together with objects such as walls, ceiling, people and furniture.

Figure 1 is a schematic illustration of an example terrain modelling scenario in which a laser scanner 2 is used to scan a terrain area 4. In this scenario, the laser scanner 2 used to scan a terrain area 4 is a Riegl laser scanner.

The laser scanner 2 generates dense 3D point cloud data for the terrain area 4 in a conventional way. This data is sent from the laser scanner 2 to a processor 3.

In this example, the terrain area 4 comprises an area of ground 6 (or terrain surface), and two objects, namely a building 8 and a tree 10.

The generated 3D point cloud data for the terrain area 4 is processed by the processor 3 using an example terrain modelling algorithm, hereinafter referred to as the "segmentation algorithm", useful for understanding the invention. The segmentation algorithm advantageously tends to provide a representation of the ground 6, as well as representations of the various objects 8, 10 above the ground 6, and also to the refinements that can be made to the representation of the ground 6 using the representations of the objects 8, 10, as described in more detail later below.

Figure 2 is a process flowchart showing certain steps of an example of a process implemented by the segmentation algorithm performed by the processor 3.

At step s2, a ground extraction process is performed on the 3D point cloud data. The ground extraction process explicitly separates 3D point cloud data corresponding to the ground 6 from that corresponding to the other objects, i.e. here the building 8, and the tree 10 and is described in more detail later below with reference to Figure 3.

At step s4 an object segmentation process is performed on the 3D point cloud data. The object segmentation process segments the 3D point cloud data such that each segment of data corresponds to a single object, as described in more detail later below with reference to Figure 5.

Figure 3 is a process flowchart showing certain steps of the ground extraction process of step s2 of the segmentation algorithm.

At step s6, a Mean Elevation Map of the terrain area 4 is computed. This is a conventional Mean Elevation Map. The resolution of a grid underlying the map may be any appropriate value.

5 In this example, the Mean Elevation Map is a grid having a plurality of cells. Each cell has assigned to it a height value determined from height values corresponding to laser sensor returns from that cell. In this example, the height value for a cell is the average of the height values corresponding to laser sensor returns from that cell.

At step s8, a surface gradient value is computed for each cell in the grid.

10 A surface gradient value for a particular cell is obtained by first computing the gradients between that cell and each of the surrounding cells. The gradient with the largest absolute value is retained as the gradient at the particular cell.

15 At step s10, cells corresponding to relatively flat surfaces are identified. In this example, this is achieved by selecting cells having a surface gradient value below a gradient-threshold value. In this example, the gradient-threshold value is 0.5. This corresponds to a slope angle of 27 degrees. However, in other examples a different gradient-threshold value is used.

20 At step s12, the cells identified as corresponding to the relatively flat surfaces, i.e. the cells that have a surface gradient value below the gradient-threshold, are grouped together with any adjacent cells having a surface gradient value below the gradient-threshold value. This forms clusters of cells that correspond to relatively flat areas.

At step s14, the largest cluster of cells that correspond to a relatively flat area, i.e. the cluster formed at step s12 containing the largest number of cells, is identified.

25 At step s16, the identified largest cluster is used as a reference cluster with respect to which it can be determined whether the other smaller clusters formed at step s12 correspond to the ground 6 of the terrain area 4. The reference cluster is used because locally smooth clusters that do not correspond to the ground 6 may exist. Thus, these cases are filtered out using the reference to the ground 6 provided by the largest ground cluster.

30 In this example, the identified largest cluster is assumed to correspond to the ground 6. Thus, any of the smaller clusters of cells, the cells of which have substantially smaller or larger height values than those of the largest cluster, are assumed not to correspond to the ground 6. In other words, in this example the cells corresponding to the ground 6 is defined to be the union of the largest cluster of cells, with a surface gradient value below the gradient-threshold, and the other clusters of cells, also containing a surface gradient value

below the gradient-threshold, in which the absolute value of the average height of the cells minus the average height of the cells in the largest cluster is smaller than a height-threshold. In this example, this height-threshold is 0.2m.

At step s18, a correction of errors generated during the computations of the surface gradient values is performed. One source of such errors and the correction of those errors will now be explained with reference to Figure 4.

Figure 4 is a schematic illustration of three cells of the grid of the Mean Elevation Map, namely the first cell 12, the second cell 14, and the third cell 16.

The height value for the first cell 12, i.e. the average of the height values corresponding to laser sensor returns from the first cell 12, is hereinafter referred to as the "first height value 18".

The height value for the second cell 14, i.e. the average of the height values corresponding to laser sensor returns from the second cell 14, is hereinafter referred to as the "second height value 20".

The height value for the third cell 16, i.e. the average of the height values corresponding to laser sensor returns from the third cell 16, is hereinafter referred to as the "third height value 22".

In this example the first height value 18 and the second height value 20 are substantially equal. Also, the third height value 22 is substantially greater than the first height value 18 and the second height value 20.

The surface gradient value for the second cell 14, which is determined at step s8 as described above, is obtained by first computing the gradients between that cell and each of the surrounding cells. The gradient with the largest absolute value is retained as the gradient at the particular cell. Thus, in this example the surface gradient value for the second cell 14 is the slope between the height levels of the second cell 14 and the third cell 16 (since the gradient between the first cell 12 and the second cell 14 is zero). This gradient is indicated in Figure 4 by the reference numeral 24. Thus, in this example the second cell has a relatively large surface gradient value. In particular, the service gradient value of the second cell 14 is above the gradient-threshold. Thus, the second cell 14 is included in the same cluster of cells as the first cell 12 despite the second height value 20 being substantially equal to the first height value 18.

Such errors are corrected at step s18 of the ground extraction process as follows. Each cell identified as not belonging to the ground is inspected. The neighbour cells of the

cell being inspected that correspond to the ground 6 are identified, and their average height is computed. If the absolute value of the difference between this average height and the height in the inspected cell is less than a correction-threshold value, the inspected cell is identified as corresponding to the ground 6. For example, returning to Figure 4, the first cell 12 corresponds to the ground 6, whereas the third cell 16 corresponds to an object 8, 10. The difference between the height of the first cell 12, i.e. the first height value 18, and the height of the second cell 14, i.e. the second height value, is zero. In this example the correction-threshold is 0.1m. Thus, since zero is less than 0.1m the second cell 14 is identified as corresponding to the ground 6.

At step s20, the steps s12 and s14 as described above are repeated. The correction of errors carried out at step s18 modifies the cluster of cells that correspond to the ground. Thus, the operations carried out steps s12 and s14, i.e. the forming clusters of cells that correspond to areas of relatively flat terrain and the identification of the largest cluster of cells, are repeated after performing the function of step s18 to accommodate for the changes made.

The correction steps s18 and s20 allow for the reconstructing of a larger portion of the ground 6 of the terrain area 4. This is because a reconstruction of the ground 6 obtained without this correction comprises a number of "holes" that are not identified as either the ground 6 or an obstacle 8, 10. The performance of the correction steps s18 and s20 advantageously tends to remove these holes. This may, for example, allow a path planner to find paths going through areas of the map previously marked as containing obstacles.

Thus, the ground extraction process of step s2 is performed. Returning to Figure 2, this ground extraction process is followed by the object segmentation process of s4.

Figure 5 is a process flowchart showing certain steps of the object segmentation process of step s4 of the segmentation algorithm.

At step s22, a Min-Max Elevation Map of the terrain area 4 is computed. This is a conventional Min-Max Elevation Map. The resolution of a grid underlying the map may be any appropriate value. In this example, the grid of the Min-Max Elevation Map is the same as that of the Mean Elevation Map. This Min-Max Elevation Map of the terrain area 4 is hereinafter referred to as the global map.

In this example, the Min-Max Elevation Map is a grid having a plurality of cells. Each cell has assigned to it a range of height values. The range of height values assigned to a particular cell ranges from the minimum to the maximum height values corresponding to laser sensor returns from that cell.

Figure 6 is a schematic illustration of the first cell 12, the second cell 14, and the third cell 16 of the grid and the range of height values assigned to each of these cells in the Mean Elevation Map.

In this example, the first cell 12 is assigned a range of height values, hereinafter referred to as the "first range 26". The first range 26 has a minimum, indicated in Figure 6 by the reference numeral 260, and a maximum, indicated in Figure 6 by the reference numeral 262.

Also, the second cell 14 is assigned a range of height values, hereinafter referred to as the "second range 28". The second range 28 has a minimum, indicated in Figure 6 by the reference numeral 280, and a maximum, indicated in Figure 6 by the reference numeral 282.

Also, the third cell 16 is assigned a range of height values, hereinafter referred to as the "third range 30". The third range 30 has a minimum, indicated in Figure 6 by the reference numeral 300, and a maximum, indicated in Figure 6 by the reference numeral 302.

Thus, the first, second, and third cells 12, 14, 16 each have a volume assigned to them that represents the range of the heights corresponding to the laser returns from that cell.

At step s24, adjacent cells corresponding to an object 8, 10, i.e. the sets of cells not identified as corresponding to the ground 6 at step s16 of the ground extraction process, are connected together to form clusters of object cells.

At step s26, for each identified object cluster a second Min-Max Elevation Map is built from the laser returns contained in that cluster. These second Min-Max Elevation Maps are hereinafter referred to as "local maps". The local maps have higher resolution than the global map generated at step s22. For example, the cell size in the local maps is 0.2m by 0.2m, whereas the cell size in the global map is 0.4m by 0.4m.

At step s28, for each local map, the range of height values of each cell in the local map is divided into segments, or voxels. Each voxel for a cell corresponds to a sub-range of the range of height values. In this example, the height of each voxel is 0.2m. However, in other examples a different voxel height is used.

Each voxel contains the height values within the corresponding sub-range of the laser returns from that cell. Voxels that do not contain any laser returns are disregarded. Also, voxels of a particular cell are merged with other voxels of that cell, if they are in contact with those other voxels.

Figure 7 is a schematic illustration of the first cell 12, the second cell 14, and the third cell 16 of the Mean Elevation Map grid, after performing step s28.

In this example, the third cell 16 was identified as corresponding to an object, i.e. not corresponding to the ground. Thus, a higher resolution grid is defined over the third cell 16, and the range of values of laser returns in each of the cells of the higher resolution grid is divided into voxels, as shown in Figure 6. In this example, each of the cells of the higher resolution grid of the third cell 16 contain the same data. Also, only the voxels corresponding to the higher height values in the third range 30 and the lower height values in the third range 30 contain any laser scanner returns. Voxels in the middle of the third range 30 do not contain any laser scanner returns. Thus, in this example, each of the cells of the higher resolution grid of the third cell 16 contains two voxels, one containing laser scanner returns corresponding to relatively lower height values, and the other containing laser scanner returns corresponding to relatively higher height values. The voxels corresponding to lower height values are hereinafter referred to as the "lower voxels" and are indicated in Figure 7 by the reference numeral 38. The voxels corresponding to higher height values are hereinafter referred to as the "upper" voxels and are indicated in Figure 7 by the reference numeral 38.

At step s30, the voxels corresponding to the ground 6 are identified. The identification of these voxels is implemented as follows. For a given cell, a number of the closest cells corresponding to the ground 6 in the grid are identified. If the absolute value of the difference between the mean height value of the lowest voxel in the given cell and the mean of the heights of the closest cells, is less than a voxel-threshold, then the that voxel is marked as corresponding to the ground 6. For example, the lowest voxels in the third cell 16 are the lower voxels 36. The second cell 14 may be identified as a closest cell that corresponds to the ground 6 to the third cell 16. In this example, the mean height of the second voxel 28 and the lower voxels 36 are substantially the same, i.e. the difference between these values is below a voxel-threshold value of, for example, 0.2m. Thus, the lower voxels 36 is identified as corresponding to the ground 6.

This process advantageously tends to allow for the reconstruction of the ground 6 under overhanging structures, for example the canopy of the tree 10.

This process also advantageously allows the reconstruction of the ground 6 that was generated at step s2 as described above with reference to Figures 2 and 3, to be refined. This is carried out at step s32.

At step s32, the reconstruction of the ground 6 is refined. At this step the fact that a voxel from a local map corresponds to the ground 6 is used to update the Mean Elevation

Map generated in the ground extraction process of s2. In particular, the cell in the Mean Elevation Map which most closely corresponds to the cell in the local map that contains the voxel corresponding to the ground 6 is identified. The identified cell is then updated by re-computing the mean height in that cell computed using only the laser returns that falls into the voxel corresponding to the ground 6.

Thus, the reconstruction of the ground 6 under overhanging structures is performed. This process advantageously exploits interaction between the Mean Elevation Map of the ground extraction process of step s2 and the Min-Max Elevation Map of the object segmentation process of step s4.

At step s34, contacting voxels are grouped together to form voxel clusters. In this example, voxels identified as belonging to the ground are interpreted as separators between clusters.

At step s36, noisy laser scanner returns are identified. In this example, voxels which contain noisy returns are assumed to satisfy the following conditions. Firstly, the voxel belongs to a cluster which is not in contact with a cell or voxel corresponding to the ground 6. Secondly, the size of the cluster (in each of the x-, y-, and z-directions) that the voxel belongs to is smaller than a predetermined noise-threshold. In this example, the noise threshold is 0.1m.

At step s38, the identified noisy returns are removed or withdrawn from the map.

This completes the segmentation algorithm performed by the processor 3.

The reconstruction of the terrain area 4 produced by performing the segmentation algorithm advantageously reconstructs portions of the ground that are under overhanging structures, for example the canopy of the tree 10. This is achieved by the steps s28 – s30 as described above.

A further advantage of the segmentation algorithm is that fine details tend to be conserved. For example, frames of windows of the building 8 are conserved by the segmentation algorithm.

The segmentation algorithm advantageously tends to benefit from the advantages of the Mean Elevation Map approach. In particular, the segmentation algorithm tends to be able to generate smooth surfaces by filtering out noisy returns.

Also, the segmentation algorithm advantageously tends to benefit from the advantages of the Min-Max Elevation Map approach. In particular, the segmentation algorithm does not make an approximation of the height corresponding to the laser scanner

return when separating the objects above the ground. Also, the local maps have higher resolution than the global map which tends to allow efficient reasoning in the ground extraction process at a lower resolution, yet provide a fine resolution object model.

A further advantage provided by the above described segmentation algorithm is that it tends to be able to achieve the following tasks. Firstly, the explicit extraction of the surface of ground 6 is performed, as opposed to extracting 3-dimensional surfaces without explicitly specifying which of those surfaces correspond to the ground 6. Secondly, overhanging structures, such as the canopy of the tree 10, are represented. Thirdly, full 3-dimensional segmentation of the objects 8, 10 is performed. Conventional algorithms do not jointly perform all of these tasks.

A further advantage of the segmentation algorithm is that errors that occur when generating 3-dimensional surfaces corresponding to the ground 6 tend to be minimised. This is due to the ability of the ground-object approach implemented by the segmentation algorithm to separate the objects above the ground.

A further advantage is that by separately classifying terrain features, the terrain model produced by performing the segmentation algorithm tends to reduce the complexity of, for example, path planning operations. Also, high-resolution terrain navigation and obstacle avoidance, particularly those obstacles with overhangs, is provided. Moreover, the segmentation algorithm tends to allow for planning operations to be performed efficiently in a reduced, i.e. 2-dimensional workspace.

Also, the provided segmentation algorithm allows a path-planner to take advantage of the segmented ground model. For example, clearance around obstacles with complex geometry can be determined. This allows for better navigation through regions with overhanging features.

In the above examples, the average value of the measured plurality of values of cells is used to determine clusters and further process those clusters, and various other processes. However, this need not be the case, and instead, in other examples, other functions may be used instead of average value, for example an average value of parameter values that remain after certain extreme values have been filtered out, or for example statistical measures other than an average as such.

In the above examples, the measured parameter is the height of the terrain and/or objects above the ground. However, this need not be the case, and in other examples any other suitable parameter may be used instead, for example colour/texture properties, optical density, reflectivity, and so on.

Apparatus, including the processor 3, for implementing the above arrangement, and performing the method steps described above, may be provided by configuring or adapting any suitable apparatus, for example one or more computers or other processing apparatus or processors, and/or providing additional modules. The apparatus may comprise a
5 computer, a network of computers, or one or more processors, for implementing instructions and using data, including instructions and data in the form of a computer program or plurality of computer programs stored in or on a machine readable storage medium such as computer memory, a computer disk, ROM, PROM etc., or any combination of these or other storage media.

10 In the above examples, the 3-dimensional point cloud data for the terrain area was provided by a Riegl laser scanner. However, in other examples the laser data is provided by a different means, for example by SICK and Velodyne sensors. Moreover, in other examples the data on the terrain area is not laser scanner data and is instead a different appropriate type of data, for example data generated by an infrared camera.

15 In the above examples, the terrain area is outdoors and comprises a building and a tree. However, in other examples the terrain area is a different appropriate area comprising any number of terrain features. In particular, the terrain features are not limited to trees and buildings.

20 In the above examples, the segmentation algorithm is performed by performing each of the above described method steps in the above provided order. However, in other examples certain method steps may be omitted. For example, steps s36 and s38 of the segmentation algorithm may be omitted, however the resulting terrain model would tend to be less accurate than if these steps were included.

25 In the above examples, the segmentation algorithm does not take into account occluded, or partially hidden, objects. However, in other examples provision is made for partially hidden objects, as will now be described in more detail with reference to Figure 8.

Figure 8 is a schematic illustration of the laser scanner 2 scanning an object that is hidden behind a further object. The object being scanned by the laser scanner is hereinafter referred to as the "hidden object 40", and the object partially hiding, or occluding, the hidden
30 object 40 is hereinafter referred to as the "non-hidden object 42".

In this example, the hidden object 40 can only be partially imaged by the laser scanner 2. Thus, a height of the hidden object observed by the laser scanner, hereinafter referred to as the "observed height 44", does not correspond to the actual object height 46.

Accurate estimation of the ground height ideally considers occlusions such as these. Thus an estimation of the ground height is preferably based on non-occluded cells. A cell can be assessed as non-occluded using a ray-tracing process.

5 In a ray-tracing process a set of cells, or a trace, is computed to best approximate a straight line joining two given cells. If any of the cells in the trace do not correspond to the ground, the end cell of the trace is occluded.

Using a ray-tracing process tends to allow for occlusions to be taken into account and reliable estimates of the ground height to be computed.

10 In order to avoid resorting to an explicit ray-tracing process and decrease the amount of computations, the following approach may be adopted. As described above, the ground is extracted by applying a threshold on the computed surface gradients. Thus, there is a "smoothness constraint" between neighbour cells identified as belonging to the ground. The terminology "smoothness constraints" is used to mean that the variation of height between two neighbour ground cells is limited. Thus, the closest ground cell to an obstacle will
15 provide a reliable local estimate of the ground height, i.e. a given ground cell is connected (via "smoothness constraints") to the rest of the ground cells which implies that this cell not only provides a local estimate of the ground height but in fact it provides a globally constrained local estimate.

20 This approach advantageously tends to avoid the use of ray-tracing while providing reliable estimates of the ground height. Standard ray-tracing techniques do not use this reasoning simply because extracting the ground is not always possible.

An embodiment of an iterative segmentation algorithm will now be described. In this embodiment, data is incorporated and the terrain model is updated as the data is collected. This advantageously tends to allow a model of the terrain to be generated in real-time, as
25 the data is collected.

Figure 9 is a process flow chart showing certain steps of an iterative segmentation algorithm.

At step s40, the laser scanner 2 generates dense 3D point cloud data for the terrain area 4 the same way as in the above described embodiments.

30 At step s42, the generated data is stored in a database. Newly generated data is added to the database as it is generated. Also, in this embodiment data may be deleted from the database, for example if it is replaced by newly generated data or it is deemed to be unnecessary or irrelevant at some point in time. In this embodiment, a policy or a filter

that encodes the definition of 'irrelevant' is utilised. For example, a filter may be used to remove data older than a certain number of seconds. Another filter may be used to filter data such that a maximum data density in a region of space is maintained e.g. if there are more than a certain number of data points per cell, the oldest data points are deleted in order to maintain a maximum density. An example policy is to discard data below a certain accuracy or quality.

At step s44, a Mean Elevation Map of the terrain area 4 is computed as described above at step s6 of Figure 3. In this embodiment, the Mean Elevation Map is a grid having a plurality of cells. Each cell has assigned to it an average height determined from the height values stored in the database corresponding to that cell. In this embodiment, the height values stored in the database are iteratively being updated as new data is generated and irrelevant data is deleted. Thus, the average heights that form the Mean Elevation Map are iteratively updated.

At step s46, a surface gradient value is computed for each cell in the grid as described above for step s8 of Figure 3. As described in more detail above, the gradient values are determined using the average heights in the Mean Elevation Map. Thus, since the average heights of the Mean Elevation Map are iteratively being updated because newly generated data is added to the database and irrelevant data is deleted from the database, the determined data values are iteratively updated.

In other embodiments, different appropriate metrics to the surface gradient values may be determined in addition to or instead of the surface gradient values. These different metric values may then be used in the determination of the ground model. For example, a value of the residual from a horizontal plane of a cell, or a plane fit metric for a cell may be used. Such metrics may be used to determine a value relating to how 'flat' a plane calculated using some or all of the data points in a cell is, e.g. the deviation of the plane from a horizontal plane. A cell may be identified as corresponding to the ground if it is suitably flat. Otherwise, the cell may be identified as corresponding to an object. These metrics may be calculated incrementally, or rapidly recalculated iteratively, as new data is provided.

At step s48, a model of the ground 6 is determined. In this embodiment the model of the ground 6 is determined by performing the following: identifying cells corresponding to relatively flat surfaces (as described above with reference to step s10 of Figure 3); forming clusters of cells that correspond to relatively flat areas (as described above with reference to step s12 of Figure 3); identifying the largest cluster of cells that correspond to a relatively flat area (as described above with reference to step s14 of Figure 3); using the identified largest cluster as a reference cluster with respect to which it can be determined whether the other smaller clusters correspond to the ground 6 of the terrain area 4 (as described above with

reference to step s16 of Figure 3); correcting of errors (as described above with reference to step s10 of Figure 3); and repeating certain of the steps to generate a ground model (as described above with reference to step s20 of Figure 3). The model of the ground 6 is iteratively updated as the gradient values (and/or any other determined metrics) are iteratively updated.

Thus, the model of the ground 6 is iteratively updated as data generated by the laser scanner 2 is added to the database, and/or as data is deleted from the database. In this embodiment, the model of the ground 6 is updated at the rate data is input or removed from the database, for example continuously.

At step s50, the ground 6 and objects 8, 10 are segmented by performing the following steps that are described in more detail above at steps s22 to s32 of Figure 5: generating a Min-Max Elevation Map of the terrain area 4 using the data contained in the database (as described above with reference to step s22 of Figure 5); forming clusters of cells corresponding to objects (as described above with reference to step s24 of Figure 5); forming local Min-Max Elevation Maps for the object clusters (as described above with reference to step s26 of Figure 5); dividing each cell in each local map into voxels (as described above with reference to step s28 of Figure 5); identifying the voxels corresponding to the ground 6 (as described above with reference to step s30 of Figure 5); and refining the reconstruction of the ground 6 (as described above with reference to step s32 of Figure 5).

In this embodiment, the Min-Max Elevation Maps are grids having a plurality of cells. Each cell has assigned to it a range of heights determined from the height values stored in the database corresponding to that cell. In this embodiment, the height values stored in the database are iteratively being updated as new data is generated and irrelevant data is deleted. Also, the model of the ground 6 is iteratively updated as data generated by the laser scanner 2 is added to the database, and/or as data is deleted from the database as described above. Thus, the segmentation of the ground and the objects, which depends on the local Min-Max elevation maps and the ground model, is iteratively updated. The segmentation of the ground and the objects is iteratively updated at a rate that depends upon the processing power of the processor 3 which performs the segmentation.

In this embodiment, a Min-Max Elevation map is generated for each iteration of the method. However, in other embodiments a Min-Max map is not generated as such for each iteration of the method. For example, in other embodiments the structure of the database is such that it corresponds to that of a Min-Max elevation map. In such cases, the database structure contains at least the information of a min-max map. Also, the direct calculation of ground cells removes the need to explicitly determine a Min-Max elevation map (and perform steps s22 to s32) at each iteration of the method. In particular, the initial process

carried out when the data arrived means it has already been accurately determined which voxels correspond to the ground and which voxels correspond to objects, including any voxels corresponding to the ground under object overhangs. Thus, there is no requirement for the explicit determination of the Min-Max elevation map (i.e. steps s22 or s26), nor for the overhang correction process (i.e. step s30 and s32). Such steps may be “combined” in the form of a relatively efficient algorithm, as described in more detail below with reference to Figure 11.

At step s52, models of the objects 8, 10 are determined. In this embodiment the models of the objects 8, 10 are determined by performing the following: forming voxel clusters (as described above with reference to step s34 of Figure 5); and identifying and removing noisy laser scanner returns (as described above with reference to steps s36 and s38 of Figure 5). The models of the objects 8, 10 are iteratively updated as the data points in the voxels are iteratively updated. The object models are iteratively updated at a rate that depends upon the processing power of the processor 3 which performs the segmentation.

This completes the iterative segmentation algorithm.

The iterative segmentation algorithm advantageously allows streaming data from the laser scanner to be incrementally processed. This tends to provide a terrain model during data collection which is updated and refined as more data is collected.

The iterative segmentation algorithm tends to be advantageous over non-iterative segmentation algorithms in which all of the data is collected before a single iteration of a process of forming a terrain model is performed.

The iterative segmentation algorithm tends to allow for real-time generation and updating of a terrain model.

Figure 10 is a schematic block diagram showing certain details of a further embodiment implementing the process of Figure 9. Figure 10 represents the process in terms of the following functional modules: a database 500, an elevation map 502, a ground model 504, a ground/object segmenter 506, and an object model 508.

Streaming input data 510, i.e. 3D point cloud data generated by the laser scanner 2, is input into the database 500.

The elevation map 502, which in this embodiment is a mean elevation map, is updated based on data that has been added to the database 500 (hereinafter referred to as “added data 512”) and/or data that has been removed from the database 500 (hereinafter referred to as “deleted data 514”).

The ground model 504, which is determined using gradient values (and/or any other determined metrics) computed from the elevation map 502 as described in more detail above at step s8 of Figure 3, is updated based on the updated elevation map 502. In particular, the ground model 504 is updated using gradient values that have been changed as a result of the streaming input data 510 (hereinafter referred to as "changed gradients 516") and/or gradient values that have been deleted (hereinafter referred to "deleted gradients 518"). The formation of the ground model 504 comprises forming cell clusters, removing certain clusters having a height above a threshold, and correcting overhangs/ground artefacts. In this embodiment, when the ground model 504 has been determined using the latest updated changed gradients 516 and deleted gradients 518, an indication 520 that the ground model 504 has been determined is generated so that the ground/object segmenter 506 may perform segmentation of the ground and objects.

Also, in this embodiment the ground model 504 uses data stored in the database 500. This is indicated in Figure 10 by the dotted arrow indicated by the reference numeral 501.

In this embodiment, the database 500 and the elevation map 502 are each updated at the rate of the streaming of the data, i.e. as data is streamed. The rate of completely updating the ground model 504 i.e. the rate and/or frequency with which an indication 520 is generated, depends on the power of the processor 3, i.e. central processing unit power.

The ground/object segmenter 506 separates the ground model 504 from those of the objects. The ground/object segmenter 506 performs this function each time an indication 520 is generated. The segmentation of the ground and the objects is updated using data stored in the database 500 (this is indicated in Figure 10 by the dotted arrow indicated by the reference numeral 503) and the ground model 504 (this is indicated in Figure 10 by the dotted arrow indicated by the reference numeral 522). The object model 508 is updated using segmented object voxels 524 that are updated by the ground/object segmenter 506 using data stored in the database 500 and the ground model 504. Generation of the segmented voxels is described in more detail above with reference to step s28 of Figure 5.

In this embodiment, the rate of the updating of the ground model 504, the rate that the ground and objects are segmented, and the rate that the object model is updated, depends on the power of the processor 3, i.e. central processing unit power.

In this embodiment, the various updated items used by and/or determined by the functional modules, i.e. the streaming input data 510 (indicated by "A1" in Figure 10), the added data 512 and the deleted data 514 (indicated by "A2" in Figure 10), the changed gradients 516 and the deleted gradients 518 (indicated by "A3" in Figure 10), the forming of

clusters, removal of certain clusters, correcting of overhangs and generation of an indication 520 (indicated by "A4" in Figure 10), the indication 520 and the access to the ground model (indicated by "A5" in Figure 10), and the updated segmented object voxels (indicated by "A6" in Figure 10), are related to each other as follows. The A1 updated items are used to
5 determine the A2 updated items. The A2 updated items are used to determine the A3 updated items. The A3 updated items are used to determine the A4 updated items. The A4 updated items are used to determine the A5 updated items. The A5 updated items are used to determine the A6 updated items.

In the above embodiment, the functional modules (i.e. the database 500, the
10 elevation map 502, the ground model 504, the ground/object segmenter 506, and the object model 508) are updated using a distinct section of code for each functional module. However, in other embodiments the functional modules may be implemented in a different appropriate way. In other embodiments, two or more functional modules may be implemented by a single block of code. For example, in a further embodiment the process of
15 updating the Min-Max Elevation Map with new data points (measured height values), updating the ground model using updated gradients values, and refining the reconstruction of the ground under overhanging objects, are performed by a single iterative process, as described below with reference to Figures 11.

Figure 11 is a process flow chart showing certain steps of a ground model updating
20 process. The process of updating the ground model will be described for a single new measurement of the height parameter. However, it will be appreciated that the process may be utilised for updating the ground model for any number of new measurements, for example by performing the process iteratively.

At step s54, a new sensor measurement within the terrain area 4 is performed. In this
25 embodiment, this new sensor measurement is a measurement of the height of the terrain.

At step s56, the voxel to which the new measurement corresponds is identified. In other words, the voxel of the terrain area 4 in which the sensor measurement is performed is identified.

At step s58, it is determined whether or not there exists an empty voxel below the
30 voxel identified at step s56.

In this embodiment, an empty voxel is defined as a voxel that contains less than a certain number of data points. In other words, a voxel is defined as empty if the number of measurements that have been made in that voxel is below a threshold value. Equivalently, in this embodiment a voxel is defined as non-empty if the number of measurements that have
35 been made in that voxel is equal to or above that threshold value.

If it is determined that there exists an empty voxel below the voxel identified at step s56, the ground model updating process proceeds to step s60. However, if it is determined that there is no existence of an empty voxel below the voxel identified at step s56, the ground model updating process proceeds to step s65.

5 At step s60, the first empty voxel directly below the voxel identified at step s56 is identified.

At step s62, it is determined whether or not there exists a non-empty voxel below the empty voxel identified at step s60.

10 If it is determined that there exists a non-empty voxel below the empty voxel identified at step s60, the ground model updating process proceeds to step s64. However, if it is determined that there is no existence of a non-empty voxel below the empty voxel identified at step s60, the ground model updating process proceeds to step s65.

At step s63, the voxel corresponding to the new measurement is identified as not corresponding to the ground.

15 This is because there exists a non-empty voxel below the voxel in which the new measurement was made, and these voxels are separated by one or more empty voxels. Thus, the voxel identified at step s56 corresponds to an overhanging structure, i.e. an object above the ground 6, and the new sensor measurement is of the object above the ground.

20 At step s65, the voxel corresponding to the new measurement is identified as corresponding to the ground.

This is because no non-empty voxels that are separated from the voxel in which the new measurement was made exist below the voxel in which the new measurement was made. Thus, the voxel in which the new measurement was made corresponds to the ground 6.

25 In this embodiment, at step s66, the average height of the ground in the cell in which the new sensor measurement (i.e. the cell containing the voxel in which the new sensor measurement was made) was made is updated using the new sensor measurement. Thus step s66 represents one possible use of the new information obtained as a result of step s65. It will be appreciated that in other embodiments the information obtained at step s65
30 may be used in other ways instead of or in addition to the use made in this embodiment at step s66.

In this embodiment the average height of the ground is updated, i.e. the reconstruction of the ground 6 is refined, in the same way as described above at step s32 of

Figure 5. In particular, the cell in the Mean Elevation Map which most closely corresponds to the cell in the local map that contains the voxel in which the new measurement was made is identified, and this identified cell is then updated by re-computing the mean height in that cell using the new measurement value as well as values of previous measurements taken in that cell. In other embodiments, a different appropriate method of updating the average height of the ground may be used, such as utilising only measurement values that are measured in the uppermost layer of voxels that correspond to the ground surface.

Following the ground model updating process, the gradient values may be recalculated to produce a 'finalised' ground model. This advantageously tends to provide that new measurements made of objects connected to the ground (but not the ground itself) are not used to update the ground model.

In the above described ground model updating process, it is assessed whether the height of the ground model in a particular voxel is updated based on height measurements made in voxels below it. However, in other embodiments a different criterion for deciding whether the height of the ground model in a particular voxel is updated may be used in addition to or instead of the above described process. For example, a process of determining whether or not a new measurement corresponds to the ground (described below with reference to Figure 12) may be advantageously incorporated into the ground model updating process of Figure 11.

Figure 12 is a process flow chart showing certain steps of a process for determining whether a new measurement corresponds to the ground. In this embodiment, the process shown in Figure 12 is performed between steps s56 and s58 of Figure 11, i.e. after performing step s56, but before performing step s58. The remaining steps of Figure 11 (i.e. steps s58 to s66) are then performed on measurements not identified as not corresponding to the ground surface (i.e. in effect identified as possibly corresponding to the ground). This advantageously tends to improve the efficiency of the ground model updating process of Figure 11.

At step s100, the number of non-empty voxels connected to the voxel to which the new measurement corresponds (i.e. the voxel identified at step s56) and in the same column as the voxel to which the new measurement corresponds (i.e. corresponding to the same cell or sub-cell) is determined. In this embodiment, one voxel is referred to as 'connected' to another voxel if there is no existence of an empty voxel between the two voxels in question.

At step s102, if the number of voxels in the column is less than three, the new measurement is identified as possibly corresponding to the ground. In other words, the new measurement is identified as possibly corresponding to the ground or is identified as not

corresponding to the ground. This completes the process of Figure 12. In this embodiment, when the new measurement is identified as possibly corresponding to the ground, the process of Figure 11 continues on to step s58; whereas if the new measurement is identified as not corresponding to the ground the process of Figure 11 is terminated i.e. there is no need to perform steps s58 to s66.

Thus, in this embodiment the new measurement is identified as possibly corresponding to the ground if the number of connected non-empty voxels in the same column as the voxel corresponding to the new measurement is less than three. Also, the new measurement is not identified as corresponding to the ground if the number of connected non-empty voxels in the same column as the voxel corresponding to the new measurement is three or more.

A column of three connected non-empty voxels may be labelled as a non-ground object (and thus measurements made in these connected voxel are not used to update the ground model) without the need to identify empty voxels below them. This is because measurement corresponding to the (two-dimensional) ground surface typically lie in a single voxel in a column. However, the ground surface may lie at the boundary between two voxels. Thus, measurements corresponding to the ground surface may lie within two voxels in a column. Having three connected voxels being non-empty tends to require that the surface being measured is at least one voxel thick. This is typically thicker than the two-dimensional ground surface, and thus a column of three or more connected non-empty voxels is assumed to correspond to an object. In this example, measurements made in columns of less than three connected non-empty voxels are processed, for example using the ground model updating process described above with reference to Figure 11.

In this embodiment, the number of connected non-empty voxels required to be in the column in order for the column to be classified as an object is three. However, in other embodiments a different number of voxels may be required.

Thus, a sufficiently thick (e.g. three voxels thick) column of connected non-empty voxels is too thick to be the two dimensional ground surface, and so must correspond to an object. There is no need to calculate the mean height of the connected voxel column and then the surrounding gradients in this case.

The process of determining whether a new measurement corresponds to the ground described above with reference to Figure 12 may then be followed by the remaining steps of the ground model updating process of Figure 11 (i.e. steps s58 to s66) which may be performed on those measurements not identified as corresponding to the ground.

In a further embodiment, the above described method of determining whether a new measurement corresponds to the ground may be performed without performing steps s58 to s66 of Figure 11 (i.e. steps s54 and s56 of Figure 11 may be performed, followed by steps s100 and s102 of Figure 12) and the updating of the ground model may be performed just using the measurements identified as corresponding to the ground in this method. Here, the new measurements not identified as corresponding to the ground may be identified as not corresponding to the ground.

The above described ground model updating process (described above with reference to Figure 11 and, optionally, Figure 12) tends to advantageously allow for the efficient updating of the ground model.

A further advantage of the above described ground model updating process is that its complexity is the same for each new measurement value, i.e. if the process is used iteratively to update the ground model with a series of new data points, the complexity of the algorithm is the same for each iteration. Also, the ground model updating process advantageously comprises a limited number of operations which are performed on floating point numbers.

In the above embodiments, the ground model updating process is used to iteratively update the ground model generated using an above described segmentation process. The ground model is updated using selected new height measurements (i.e. those measurements that only correspond to the ground). However, in other embodiments the ground model updating process may be used to update a ground model that is generated using a different appropriate method. For example, the ground model updating process may be used to update ground models generated using the Mean Elevation Map, the Min-Max Elevation Map, the Multi-Level Elevation Map, the Volumetric Density Map, Ground Modelling via Plane Extraction, and Surface Based Segmentation. Thus it will be appreciated that the processes described with reference to Figure 11 in themselves provide embodiments of the present invention.

In other words, the ground model updating process for updating the average height value of the ground in a cell or voxel with a new height measurement (which includes identifying the voxel within which the new height measurement lies, and updating the average height of the ground only if the voxel within which the measured parameter value lies coincides the voxels corresponding to the ground surface or the voxel within which the measured parameter value lies and voxels corresponding to the ground surface are not separated by non-empty voxels) may be used to update a ground model generated by any appropriate process.

The above embodiments of a segmentation algorithm provide a model of the ground surface and three dimensional models of objects above the ground. In further embodiments one or more of the object models that have been generated are then classified, i.e. the object model is assigned a particular object-class. Thus, the object corresponding to the object model is identified as a particular type of object. Classification processes that may be used for the classifying of the object model include those classification processes that compare features of the object model to features of one or more template models representative of one or more respective object classes. For example, conventional feature based classification process that incorporate Principal Component Analysis (PCA), a 'Spin Image', Moments Grids, and/or Spherical Harmonic Descriptors may be used.

Classification of the object models generated by the above described segmentation algorithms tends to produce more accurate and efficient results compared to classification of object models generated using conventional techniques. One reason for this is that better segmentation of the objects and the ground tends to result from the above described segmentation algorithm.

In the claims which follow and in the preceding description of the invention, except where the context requires otherwise due to express language or necessary implication, the word "comprise" or variations such as "comprises" or "comprising" is used in an inclusive sense, i.e. to specify the presence of the stated features but not to preclude the presence or addition of further features in various embodiments of the invention.

It is to be understood that, if any prior art publication is referred to herein, such reference does not constitute an admission that the publication forms a part of the common general knowledge in the art, in Australia or any other country.

CLAIMS:

1. A process of identifying a value of a parameter measured in a cell in an extraction process for extracting an object (8, 10) or terrain feature (6) as corresponding to a particular object (8, 10) or terrain feature (6), the process comprising:

5 for a measurement of the parameter in the cell:

dividing a range of possible parameter values into a plurality of voxels such that the value of the measurement of the parameter in the cell lies within one of the voxels; and

identifying the measured value of the parameter as corresponding to the particular object (8, 10) or terrain feature (6) if:

10 the voxel within which the measured parameter value lies coincides with a range of values that corresponds to a particular object (8,10) or terrain feature (6); or

there is no voxel between the voxel within which the measured parameter value lies and a voxel that coincides with a range of values that corresponds to a particular object (8,10) or terrain feature (6) within which below a certain number of measurements of the parameter have been made.

2. A process according to claim 1, further comprising:

for a measurement of the parameter in the cell:

identifying the measured value of the parameter as possibly corresponding to the particular object (8, 10) or terrain feature (6) if:

20 there does not exist two different voxels which are different from the voxel within which the measured parameter value lies such that there are there is no voxel between the voxel within which the measured parameter value lies and each of the two different voxels within which below a certain number of measurements of the parameter have been made.

3. A process according to claim 1 or 2, further comprising updating a value of a function only if the value of the measured parameter is identified as corresponding to the particular object (8, 10) or terrain feature (6).

4. An apparatus for generating a model of terrain in accordance with the process of any of claims 1 to 3, the apparatus comprising scanning and measuring apparatus (2) for measuring values of a parameter, and one or more processors (3) arranged to perform the processing steps recited in claims 1 to 3.

5. A computer program or plurality of computer programs arranged such that when executed by a computer system it/they cause the computer system to operate in accordance with the process of any of claims 1 to 3.

6. A machine readable storage medium storing a computer program or at least one of the plurality of computer programs according to claim 5.

1/10




FIG. 1

2/10




FIG. 2

3/10




FIG. 3




FIG. 4




FIG. 8

5/10




FIG. 5

6/10




FIG. 6




FIG. 7

7/10




FIG. 9

8/10






FIG. 10

9/10



10/10

**FIG. 12**

INTERNATIONAL SEARCH REPORT

International application No.

PCT/AU2011/000017

A. CLASSIFICATION OF SUBJECT MATTER

Int. Cl.

G06T 17/05 (2011.01)

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

WPI & EPODOC: Keywords- object, terrain, land, geo, voxel, 3d segment, parameter, height, elevation and like words.

Google Patents, Patent Lens and Google Scholar: Similar keywords as above.

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
P, X	DOUILLARD et al., 'Hybrid Elevation Maps: 3D Surface Models for Segmentation', The 2010 IEEE/RSJ International Conference on Intelligent Robots and Systems, 18 October 2010, Taipei, Taiwan; pages 1532-1538. See whole document.	1 and 3-6
X	US 7,272,474 B1 (STENTZ et al.) 18 September 2007 See figure 10; column 8, lines 48-65.	1 and 3-6
A	SUZUKI et al., 'A Pattern Matching Technique for Detecting Similar 3D Terrain Segments', AIKED'05 Proceedings of the 4th WSEAS International Conference on Artificial Intelligence, Knowledge Engineering Data Bases; 2005. See whole document.	

☐ Further documents are listed in the continuation of Box C☒ See patent family annex

* Special categories of cited documents:	
"A" document defining the general state of the art which is not considered to be of particular relevance	"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
"E" earlier application or patent but published on or after the international filing date	"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)	"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art
"O" document referring to an oral disclosure, use, exhibition or other means	"&" document member of the same patent family
"P" document published prior to the international filing date but later than the priority date claimed	

Date of the actual completion of the international search
07 March 2011Date of mailing of the international search report
10 MAR 2011

Name and mailing address of the ISA/AU

AUSTRALIAN PATENT OFFICE
PO BOX 200, WODEN ACT 2606, AUSTRALIA
E-mail address: pct@ipaustalia.gov.au
Facsimile No. +61 2 6283 7999

Authorized officer

DEBASISH DASGUPTA
AUSTRALIAN PATENT OFFICE
(ISO 9001 Quality Certified Service)
Telephone No : +61 2 6283 2587

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No.

PCT/AU2011/000017

This Annex lists the known "A" publication level patent family members relating to the patent documents cited in the above-mentioned international search report. The Australian Patent Office is in no way liable for these particulars which are merely given for the purpose of information.

Patent Document Cited in Search Report	Patent Family Member
US 7272474	NONE
Due to data integration issues this family listing may not include 10 digit Australian applications filed since May 2001.	
END OF ANNEX	