



(43) International Publication Date
20 September 2012 (20.09.2012)

- (51) International Patent Classification:
G06T 17/05 (2011.01) *G06K 9/46* (2006.01)
- (21) International Application Number:
PCT/AU2012/000248
- (22) International Filing Date:
9 March 2012 (09.03.2012)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:
2011900895 11 March 2011 (11.03.2011) AU
- (71) Applicant (for all designated States except US): **THE UNIVERSITY OF SYDNEY** [AU/AU]; Sydney, New South Wales 2006 (AU).
- (72) Inventors; and
- (75) Inventors/Applicants (for US only): **UNDERWOOD, James Patrick** [AU/AU]; Unit 9/111 McEvoy Street, Alexandria, New South Wales 2015 (AU). **DOUILLARD, Bertrand** [FR/AU]; 90A Shepherd Street, Chippendale, New South Wales 2009 (AU). **MORTON, Peter Michael** [AU/AU]; 61 Marina Road, Mentone, Victoria 3194 (AU).
- (74) Agents: **NEGLER, Justin Pierre** et al.; DAVIES COLLISON CAVE, 1 Nicholson Street, Melbourne, Victoria 3000 (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, QA, RO, RS, RU, RW, 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, RW, 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



(57) Abstract: A method and apparatus for processing an image, the method comprising: generating a plurality of data points by measuring a value of a parameter at a plurality of different locations in an area of terrain (4) (e.g. by using a laser scanner); performing a segmentation process using the plurality of data points to generate a plurality of image segments, each image segment corresponding to an object (8, 10) or terrain feature (6) (for example, the ground); selecting data points of a particular one of the generated plurality of image segments; and performing a tracking algorithm using the selected data points to track a particular object (8, 10) or terrain feature (6) that corresponds to the particular image segment. (2). The measured parameter may be a three dimensional position.

- 1 -

EXTRACTION PROCESSES

FIELD OF THE INVENTION

The present invention relates to extraction and/or segmentation, extraction and/or segmentation processes, extraction and/or segmentation algorithms, and
5 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
10 laser scanner, such as a Velodyne 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
15 terrain area for use in, for example, path planning or analysing mining environments.

Tracking algorithms that use 3D point cloud data of an area, to track objects in that area have been developed.

However, conventional approaches typically project the 3D point cloud into
20 a 2D representation. For example, the Stanford entry into the 2007 DARPA Urban Challenge comprised a method in which Velodyne data was re-sampled into 2D "virtual scans" in order to perform segmentation and tracking.

By projecting the 3D point cloud into a 2D representation, conventional tracking algorithms tend not to exploit all the information available.

25

SUMMARY OF THE INVENTION

In a first aspect, the present invention provides a method for processing an image, the method comprising: generating a plurality of data points by measuring

- 2 -

a value of a parameter at a plurality of different locations in an area of terrain; performing a segmentation process using the plurality of data points to generate a plurality of image segments, each image segment corresponding to an object or terrain feature; selecting data points of a particular one of the generated plurality of
5 image segments; and performing a tracking algorithm using the selected data points to track a particular object or terrain feature that corresponds to the particular image segment.

The measured parameter may be a three dimensional position.

The tracking algorithm may comprise filtering the selected data points.

10 The tracking algorithm may comprise, for each selected data point, determining a vector that is normal to a surface of the image segment that corresponds to the particular object or terrain feature at that data point.

The tracking algorithm may comprise, for each selected data point, determining a range based descriptor for that data point.

15 The tracking algorithm may further comprise, for each selected data point, determining a coordinate vector.

A coordinate vector for a selected data point may comprise: a value or values corresponding to a time at which that data point was measured; a value or values corresponding to the position measured for that data point; the range based
20 descriptor determined for that data point; the surface normal determined for that data point; and an identity value for image segment that that data point belongs to.

The tracking algorithm may comprise determining a centroid of the selected data points, and tracking the determined centroid.

The segmentation algorithm may comprise: determining a terrain mesh
25 comprising nodes connected by edges, the nodes corresponding to the data points; for each node, determining a value of a function of the measured parameter values; selecting at least one node that corresponds to a determined function value that is below a threshold value; identifying the selected at least one node as corresponding to the particular object or terrain feature; and identifying a

- 3 -

node other than the selected at least one node as corresponding to the particular object or terrain feature if: that node corresponds to a determined function value that is below the threshold value; and that node and the selected at least one node are connected in the terrain mesh by a route, the route comprising either: a single
5 edge; or a plurality of edges and one or more nodes, wherein each node on the route corresponds to a determined function value that is below the threshold value.

The step of identifying a node other than the selected at least one node as corresponding to the particular object or terrain feature may comprise: identifying the selected at least one node as the current node; and iteratively performing
10 steps (i) to (iii); wherein: step (i) comprises identifying a node connected in the terrain mesh to the current node by a single edge; step (ii) comprises identifying the node identified at step (i) as corresponding to the particular object or terrain feature if that node corresponds to a determined function value that is below the threshold value; and step (iii) comprises identifying the node identified at step (i)
15 as the current node if, at step (ii), the node identified at step (i) is identified as corresponding to the particular object or terrain feature.

The step of, for each node, determining a value of a function using the measured parameter values may comprise: for each edge connected to that node: identifying a node that is at a distance greater than a threshold distance from the
20 node in question in a direction along that edge; and determining a value of a gradient between the node in question and the identified node; and determining a value of a function of the gradient values.

The function of the gradient values may be a maximum of the absolute values of the gradient values.

25 The method may further comprise identifying as corresponding to a different object or terrain feature to the particular object or terrain feature, each of the nodes in the terrain mesh not identified as corresponding to the particular object or terrain feature.

In a further aspect, the present invention provides apparatus for processing
30 an image, the apparatus comprising: a sensor arranged to measure a value of a

- 4 -

parameter at a plurality of different locations in an area of terrain so as to generate a plurality of data points; and one or more processors arranged to: perform a segmentation process using the plurality of data points to generate a plurality of image segments, each image segment corresponding to an object or terrain feature; and perform a tracking algorithm using data points of one of the generated plurality of image segments so as to track a particular object or terrain feature that corresponds to that image segment.

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 any of the above aspects.

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.

15

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 an embodiment of a segmentation and tracking algorithm;

20

Figure 3 is a process flowchart showing certain steps of a segmentation algorithm;

Figure 4 is a schematic illustration (not to scale) of an example terrain mesh constructed for the terrain area using range data from the laser scanner;

25

Figure 5 is a process flowchart showing certain steps of the ground extraction process of the segmentation algorithm;

Figure 6 is a process flowchart showing certain steps of a tracking algorithm;

- 5 -

Figure 7 is a process flow chart showing certain steps of an example of a method for evaluating the performance of a segmentation algorithm;

Figure 8 is a schematic illustration (not to scale) of a manually determined segmentation of a tree; and

5 Figure 9 is a schematic illustration (not to scale) of a segmentation of the tree determined using the segmentation algorithm.

DETAILED DESCRIPTION

The terminology "terrain" and "terrain features" are used herein to refer to a
10 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
15 setting, either indoors or outdoors. The geometric configuration of other objects or structures above this surface, may, for example, include 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,
20 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
25 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 Velodyne laser scanner. The laser scanner 2 may be positioned at any appropriate position relative to the
30 terrain area 4. For example, the laser scanner 2 may be positioned above the

- 6 -

terrain area 4 looking down, or the laser scanner may be positioned on the terrain etc.

The laser scanner 2 generates 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.

5 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 (i.e. the range image, or image taken by the laser scanner 2) is processed by the processor 3 using a segmentation and tracking algorithm, an embodiment of which is
10 described in more detail later below with reference to Figure 2.

Figure 2 is a process flowchart showing certain steps of an embodiment of a segmentation and tracking algorithm performed by the processor 3.

At step s100, a segmentation algorithm is performed on the 3D point cloud data to extract segments corresponding to the objects 8, 10.

15 In this embodiment, the segmentation algorithm is a mesh-based segmentation algorithm. The segmentation algorithm used in this embodiment is described in more details later below with reference to Figure 3 to 5. However, in other embodiments, a different appropriate segmentation algorithm is used to extract the object segments.

20 At step s200, a tracking algorithm is performed on each of the extracted object segments.

In this embodiment, each of the object segments are tracked using the tracking algorithm. However, in other embodiments, a sub-set of the object segments are tracked.

25 The tracking algorithm used in this embodiment will be described in more detail later below after the description of the segmentation algorithm and with reference to Figure 6.

Thus, a segmentation algorithm is provided.

- 7 -

The mesh-based segmentation algorithm tends to provide an explicit differentiation between ground and non-ground points. This explicit differentiation means that, during the performance of the tracking algorithm, ground points can be eliminated from further processing. This tends to greatly reduce the computational load on the processor 3 as described in more detail later below.

A further advantage provided by the tracking and segmentation algorithm is that the full 3D point cloud data is exploited in order to track an object. This is in contrast to conventional techniques which typically project 3D point cloud data into a 2D representation.

10 A further advantage provided by the tracking and segmentation algorithm is that objects are not assumed to be of any particular shape or size. This tends to be in contrast to conventional method in which objects are assumed to be of a particular size/shape, for example, the DARPA Urban Challenge trackers typically assumed the only moving entities were cars.

15 A further advantage provided by the tracking and segmentation algorithm is that only weak assumptions about the movement of an object are made (e.g. constant velocity but with a high amount of uncertainty). Therefore, the provided segmentation and tracking algorithm tends to be able to be used to track objects belonging to different object classes (i.e. without any knowledge of what they are).
20 For example, using the above provided segmentation and tracking algorithm, a person is tracked no differently to a car, a robot, a bicycle, etc.

What will now be described is the segmentation algorithm (performed at step s100 of the segmentation and tracking algorithm of Figure 2) used in this embodiment.

25 Figure 3 is a process flowchart showing certain steps of the segmentation algorithm performed by the processor 3.

At step s1, a terrain mesh is constructed from the measured range image.

In this embodiment, the terrain mesh is constructed using the procedure disclosed in "*Segmentation of 3D Lidar Data in non-flat Urban Environments using a Local Convexity Criterion*", F. Moosmann, O. Pink, and C. Stiller, IEEE Intelligent
30

- 8 -

Vehicles Symposium, pages 215–220, 2009, which is disclosed herein by reference.

This procedure for constructing a terrain mesh exploits the raster scan arrangement of the data in a range image to generate a mesh topology. This procedure advantageously tends to be applicable to any type of range images, e.g. stereo frames for instance.

The terminology "raster scan" is used herein to refer to a rectangular pattern of image capture and reconstruction in television. The term is used herein to refer to raster graphics, i.e. the pattern of image storage and transmission used in most computer bitmap image systems. In this embodiment, in a raster scan, an image is subdivided into a sequence of (e.g. horizontal) strips known as "scan lines".

Figure 4 is a schematic illustration (not to scale) of an example terrain mesh constructed for the terrain area 4 using range data from the laser scanner 2.

In this embodiment, the terrain mesh 12 is built by first connecting neighbour returns (i.e. neighbouring data points) in each raster line 14. In this embodiment, the raster lines 14 appear as concentric circles. In this embodiment, the centre 16 of the concentric circles formed by the raster lines 14 is directly under the laser scanner 2. In this embodiment, the links between raster lines 14 are instantiated by connecting data points with the closest yaw angles. The data points form nodes of the terrain mesh 12 (three such nodes are indicated in Figure 4 by the reference numeral 18), with the raster lines 14 and links between the raster lines connecting the nodes 18.

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. the building 8, and the tree 10, and is described in more detail later below with reference to Figure 5.

At step s4 an object segmentation process is performed on the 3D point cloud data that corresponds to the objects.

- 9 -

In this embodiment, non-ground data points, i.e. the data points not identified as corresponding to the ground 6 at step s2, are connected together to form clusters of object cells.

In this embodiment, this process is performed using a conventional
5 clustering technique comprising feeding non-ground data points to a 3D grid. Voxels (the 3D cells of the grid) are then gathered to form a cluster if those cells are non-empty and in contact with each other. In other words, this process comprises assigning to a voxel each non-ground data point. The voxels are then clustered together to form "object clusters" if they are sufficiently close to one
10 another (e.g. within a specified neighbourhood). The ground operates as a separator between the object clusters. This is because, in this embodiment, data points corresponding to the ground are removed prior to voxel clustering.

In this embodiment, the non-ground data points are assigned into voxels using a voxel grid. This voxel grid may be of any size. Smaller voxels tend to allow
15 more complex representations of objects that tend to be more accurate. However, more memory is required and processing times are generally longer than if a larger voxel size is used. Furthermore, smaller voxels may lead to over segmentation because the intrinsic connectivity of objects may not be captured due to the increased number of empty voxels.

In this embodiment, the neighbourhood that is used to cluster voxels may
20 be of any appropriate size. For example, a neighbourhood may be based on the height of a given voxel above the ground 6. The use of larger neighbourhoods at relatively larger distances above the ground tends to alleviate the problem of over-segmentation of high, sparse objects while not risking under-segmentation of
25 objects that are close together and on the ground.

Optionally, the process of step s4 may require that additional criteria be satisfied. For example, in other embodiments a variance criterion must be satisfied. An example variance criterion requires that for two voxels to be clustered together, then the x, y, or z variances of the two voxels must be within a given
30 threshold of each other. This advantageously tends to separate dissimilar objects into different clusters

- 10 -

Thus, a segmentation algorithm is provided. The segmentation algorithm advantageously segments data points into a ground partition, and separate object partitions.

In this embodiment, the terrain mesh 12 is used for ground extraction, and
5 the rest of the segmentation (i.e. object segmentation) is performed in the voxel grid. This is in contrast to some conventional techniques in which full segmentation of a scene (i.e. partitioning of ground and objects) is performed using a mesh model. For example, this is in contrast to those techniques disclosed in the
10 aforementioned "*Segmentation of 3D Lidar Data in non-flat Urban Environments using a Local Convexity Criterion*", and those used in many conventional graph based segmentation methods.

Figure 5 is a process flowchart showing certain steps of the ground extraction process of step s2 of the segmentation algorithm.

At step s6, a gradient field for the terrain mesh 12 is computed.

15 In this embodiment, for each node 18 in the terrain mesh 12 a gradient is evaluated for each incoming edge connected to that node.

A so-called final gradient value for a node 18 is then determined as the maximum of the absolute values of the edge gradients determined for that node 18.

20 In other words, in this embodiment each node 18 of the terrain mesh 12 has four edges/links. These edges may be for example, identified as UP, DOWN, LEFT, and RIGHT. The UP edge from a node 18 is the edge that extends from the node 18 in a direction away from the centre 16. The DOWN edge from a node 18 is the edge that extends from the node 18 in a direction towards the centre 16.
25 The LEFT edge from a node 18 is the edge that extends from the node 18 in an anti-clockwise direction. The RIGHT edge from a node 18 is the edge that extends from the node 18 in a clockwise direction.

In this embodiment, to compute, for example, the gradient for a node 18 along its DOWN edge, the algorithm traverses the terrain mesh 12 from the node
30 18 in question to the next node in the DOWN direction (once a predetermined

minimum distance has been traversed). The gradient is then determined as the absolute difference of the height of the starting and the ending node divided by the distance between them. In this embodiment, if the predetermined distance is not traversed in a given direction, the determined gradient is disregarded.

5 The use of a predetermined minimum distance when determining gradient values for nodes 18 advantageously tends to reduce the effects of noise. In this embodiment, the nodes the closer to the centre 16 of the terrain mesh 12 are closer together than nodes 18 further from the centre 16. Thus, if consecutive nodes were used for gradients evaluation (i.e. if there were no minimum
10 predetermined distance) the determined gradient values would tend to be noisy. This may, for example, be due to the fact that the distance separating consecutive nodes is of the same order of magnitude of the noise in their sensed height.

At step s8, the main ground partition in the terrain mesh 12 is identified.

15 In this embodiment, the main ground partition is obtained by growing a cluster of ground nodes from the raster line closest to the centre 16 (i.e. the inner most circle in Figure 4) to the raster line furthest from the centre 16. This is performed as follows.

Firstly, on the closest raster line to the centre 16, a label "GROUND" is assigned to each node in the longest sequence of (adjacent) nodes that each have
20 a final gradient value below a predetermined threshold (G).

Other nodes on that raster line 14 are then assigned the GROUND label if:

- they are in a further sequence of (adjacent) nodes; and
- each of the nodes of the further sequence has a final gradient value below the predetermined threshold G; and
- 25 - the respective heights of the node in longest sequence and the node in the further sequence that are closest to one another are within a predetermined range (R).

In other embodiments, the GROUND label may be assigned to the nodes of the first raster line in a different way. For example, each node 18 on the raster line

- 12 -

closest to the centre 16 may be identified as GROUND if its final gradient value has a norm below a predetermined gradient threshold value.

Then, for the other raster lines 14 (working outwards from the raster line closest to the centre 16), a node 18 is identified as GROUND if its final gradient value has a norm below the predetermined gradient threshold value G, and at least one direct neighbour of it has already been identified as GROUND. In this embodiment, a direct neighbour of a node 18 can be a node 18 in the same raster line or in the previous raster line.

In this embodiment, the label GROUND is propagated twice in each scan line: once along each direction (i.e. each node is traversed twice).

Thus, the iterative growing of the ground partition is performed.

This process of identifying the main ground partition advantageously tends to avoid the identification of non-ground flat partitions (such as the roof of the house 8) being mistaken as the ground 6.

At step s10, a smoothing process is performed to avoid ground points being identified as part of objects 8, 10 at the junction between ground and non-ground sections.

In this embodiment, transitions between the ground 6 and the objects 8, 10 are detected as follows.

For a particular raster line, a number of consecutive nodes that have been identified as GROUND are selected. In this embodiment, ten such nodes are selected. Thus, a so-called "buffer" is selected.

In other embodiments, a different number of nodes may be used to form the buffer, e.g. five nodes.

The mean height (H) of the nodes in the buffer is determined.

The buffer is then "slid", or moved around the raster line, towards the next node. Also, the buffer is slid along each radial line of the terrain mesh 12.

In this embodiment, the next node that the buffer is slid onto is marked with the label "TRANSITION" if it does not belong to the ground 6 based on its gradient

- 13 -

(i.e. its final gradient value is above the predetermined threshold G), but its height is within the predefined range value (R) of the buffer mean height (H).

This process is repeated until all nodes 18 have been processed as described. In this embodiment, the buffer is emptied when moving on to the processing of the next scan or radial line.

In this embodiment, the buffer is slid twice along each raster line (once in a clockwise direction, and once in an anti-clockwise direction). Also, the buffer is slid twice along each radial line (once moving away from the centre 16, and once moving back towards the centre 16). This advantageously tends to identify transitions on all sides of an object.

The process of step s10 advantageously tends to allow for a tighter segmentation of objects. This can be critical when modelling particularly cluttered areas of terrain.

A further advantage of the process of step s10 is that noise/bias in the elevation data tends to be smoothed

Thus, a ground extraction process of the segmentation algorithm is provided.

The identification of the ground is an integrated part of the segmentation process. This is in contrast to other methods, in which the ground is (typically) obtained as the output of a classifier applied after segmentation.

Advantages provided by the above described segmentation algorithm is that the effects of noise on the shorter links of the terrain mesh 12 tend to be reduced. The averaging out of such noise tends to be provided for by the use of an adaptive length for the computation of the gradient values used to decide whether a data point corresponds to the ground. Such an adaptive mechanism is not provided by conventional algorithms.

The above described segmentation algorithm advantageously tends to allow for the reconstruction of the ground 6 under overhanging structures, for example the canopy of the tree 10.

- 14 -

Thus, the provided segmentation algorithm advantageously 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, 5 overhanging structures, such as the canopy of the tree 10, are represented. Thirdly, 3-dimensional segmentation of the objects 8, 10 is performed.

In this embodiment, full 3D segmentation of the data is obtained once the 3D clustering (step s4) of NON-GROUND data points (i.e. data points that are not identified as either GROUND or TRANSITION) has been performed.

10 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 15 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 20 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.

25 Furthermore, the provided segmentation algorithm advantageously tends to provide that the separation of object in the data facilitates the classification of the objects. Conventional approaches tend to attempt to classify data points without prior knowledge of their membership to a larger segment.

- 15 -

What will now be described is the tracking algorithm (performed at step s200 of the segmentation and tracking algorithm of Figure 2) used in this embodiment.

In this embodiment, the tracking algorithm is a combination of three main elements, namely object modelling, motion estimation, and data association. These take as their input a segmented 3D point cloud corresponding to an object (i.e. an object segmentation) obtained by the segmentation algorithm performed at step s100 and described in more detail above with reference to Figures 3 to 5.

Figure 6 is a process flowchart showing certain steps of a tracking algorithm in this embodiment at step s200.

At step s29, points which were labelled as "ground" during the segmentation algorithm are discarded. This elimination of the ground points from further processing advantageously tends to greatly reduce computational load.

At step s30 after segmentation, the data corresponding to each of the object segments is filtered.

In this embodiment, the data is filtered through a 870 x 64 pixel range image. Further information about this data filtering can be found in "*Segmentation of 3d Lidar data in non-flat urban environments using a local convexity criterion*", F. Moosmann, O. Pink, and C. Stiller, pages 215 –220, June 2009, which is incorporated herein by reference. This tends to re-sample the data to an equal resolution in both elevation and azimuth, and tends to allow for easier computation of surface normals and a range-based local descriptor.

At step s32, for each remaining point, a surface normal is computed (i.e. for each point, a normal to the surface of the object segment to which that point belongs is determined). Further information on the computation of surface normals can be found in "*Motion estimation from range images in dynamic outdoor scenes*", Frank Moosmann and Thierry Fraichard. In Robotics and Automation (ICRA), 2010 IEEE International Conference on, pages 142 –147, 3-7 2010, which is incorporated herein by reference.

- 16 -

At step s34, for each remaining point, a range-based local descriptor is computed. Further information on the computation of range-based local descriptors can also be found in "*Motion estimation from range images in dynamic outdoor scenes*",

5 At step s36, the filtered object points are projected into a global coordinate frame.

The projecting of Velodyne laser scanner returns into a global co-ordinate frame in real time tends to be a computationally expensive task. However, typically 40% to 70% of all Velodyne returns are classified as "ground". Thus, the
10 elimination of the ground-labelled points from the set of points that are processed when tracking is performed tends to provide that tracking can be performed in real-time.

In this embodiment, this projection of the filtered object points yields at set of projected points. Each projected point in the set is represented by a 13-
15 dimensional vector having the following form:

$$(t, x, y, z, range, descriptor, normal, segment.id)$$

where: t is a timestamp (i.e. the time at which the point was measured by
20 the laser scanner 2). In this embodiment, there is one component in each 13-dimensional vector that indicates the timestamp for a projected point;

x , y , and z represent the global position of the point in the global coordinate frame. In this embodiment, there are three components in each 13-dimensional vector that indicate the global position of a projected point
25 (x , y , and z);

$range$ is the range or distance of the point from the laser scanner 2. In this embodiment, there is one component in each 13-dimensional vector that indicate the measured range a projected point.

- 17 -

descriptor is the range based descriptor for the point determined at step s34. In this embodiment, there are four components in each 13-dimensional vector that correspond to the range based descriptor of a projected point;

5 *normal* is the surface normal for the point determined at step s32. In this embodiment, there are three components in each 13-dimensional vector that correspond to the surface normal at a projected point; and

segment.id is the output of the segmentation algorithm for that point (i.e. the classification for the point). In this embodiment, there is one
10 component in each 13-dimensional vector that indicate the classification of a projected point.

The use of this 13-dimensional vector tends to be provide that the algorithm can be performed particularly efficiently. However, in other embodiments, each projected point is represented by a different appropriate vector.

15 At step s38, each object segment is tracked using the vectors for that object segment (determined at step s36 above).

In this embodiment, any appropriate process for tracking an object segment may be used.

For example, a relatively simple method is as follows. Firstly, for each
20 object segment, the centroid of that segment (i.e. the mean of the segmented object's coordinates) is determined. A conventional process of tracking point observations is then used. An example of such a process can be found in "*Tracking methods in a multi-target environment*", Y. Bar-Shalom, Automatic Control, IEEE Transactions on, 23(4):618 – 626, August 1978, which is
25 incorporated herein by reference.

In some embodiments, a method in which all observations are treated as exact may be used. Data association may be carried out using, for example, a "nearest neighbour" process. In other embodiments, Independent Kalman Filters are applied to each of the tracks. Also, the position of subsequent observations

may be predicted. Further examples of processes of tracking point observations can be found in "*Probabilistic data association techniques for target tracking in clutter*", T. Klubarajan and Y. Bar-Shalom, Proceedings of the IEEE 92(3):536 – 557, March 2004, which is incorporated herein by reference.

5 Further examples of appropriate process for tracking an object segment use all points belonging to a segmented object as an observation (as opposed to just the centroid of the object segment), and use this as a track's appearance model. Further information on such a process may be found in "*Motion estimation from range images in dynamic outdoor scenes*".

10 An example approach that uses point clouds as observations involves creating a track with a static appearance model (e.g. those points corresponding a first observation). The track's model is then registered with subsequent observations using Point-Plane Iterative Closest Point (ICP) process. Further information about the Point-Plane ICP process can be found, for example, in
15 "*Object modelling by registration of multiple range images*", Y. Chen and G. Medioni, In Robotics and Automation, 1991, Proceedings, 1991, IEEE International Conference on pages 2724 –2729 vol.3, April 1991, which is incorporated herein by reference. A result of the Point-Plane ICP process is the determination of the object's 6 degrees of freedom displacement. This ICP tends
20 to require no explicit data association. Some form of data association may be performed however as segmented objects that do not get matched to an existing track must be instantiated as new tracks.

A matching process which exploits the results of the ICP process may be performed as follows. Firstly, as part of the ICP algorithm, links are formed
25 between the pairs of point clouds being registered. Using these links between the track's model and the segmented point cloud, the *segment.id* of each matched point is then determined. A single segment is determined using a majority vote over these *segment.id* values and gives a single segment. This segment is then associated with the track.

In other embodiments a track's appearance model is replaced with a new observation at each iteration. This tends to ensure that the appearance difference between the track and the actual object is relatively small. Thus, a good ICP match tends to be possible.

5 A further example of a process for tracking an object segment is a so-called "dynamic mapping" process. This process is described in greater detail in *"Motion estimation from range images in dynamic outdoor scenes"*. In this approach, points from each observation are accumulated into the appearance model. Points are accepted or rejected based on their measurement uncertainty so as to avoid
10 accumulating an infinite number of points. An advantage of this method is that, over time, a more complete representation of the object being tracked is formed.

The above mentioned appearance modelling techniques can be used in tracks with or without Kalman Filter estimation.

15 A further example of a process for tracking an object segment comprises a process of handling scene under- and over-segmentation. Scene under- or over-segmentation may, for example, occur when two objects approaching one another merge into one object in the segmentation result. Further details regarding an example of such an approach can be found in *"Tracking people with a 360-degree lidar"*, John Shackleton, Brian VanVoorst, and Joel Hesck, Advanced Video and
20 Signal Based Surveillance, IEEE Conference on 0:420-426, 2010, which is incorporated herein by reference.

Thus, a process of segmentation (to produce an object segment from an image) and tracking (to track that object segment through a sequence of images) is provided.

25 Apparatus, including the processor 3, for implementing the arrangements described herein, and performing the method steps described herein, 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 computer, a network of
30 computers, or one or more processors, for implementing instructions and using

- 20 -

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.

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

 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
15 terrain features are not limited to trees and buildings.

 In the above embodiments, the main ground partition in the terrain mesh 12 is identified as described above at step s8. Firstly, on the closest raster line to the centre 16, a label "GROUND" is assigned to certain nodes. Then, for the other raster lines 14, a node 18 is identified as GROUND if its final gradient value has a
20 norm below the predetermined gradient threshold value G, and at least one direct neighbour of it has already been identified as GROUND. However, in other embodiments, nodes in the other raster lines are identified if one or more different criteria are met, instead of or in addition to the criteria of the node's final gradient value having a norm below the predetermined gradient threshold value G, and at
25 least one direct neighbour of the node having already been identified as GROUND.

 For example, in other embodiments, in addition to any of the above mentioned criteria, nodes in a given raster line will only be labelled as GROUND if the previous raster line (i.e. the directly adjacent raster line that is closer to the
30 mesh centre) contains at least one node that has been identified as GROUND. This process of "growing" the ground partition iteratively from the image's centre

outwards is advantageously stopped as soon as no more ground nodes are identified in the current scan line. This technique tends to avoid attempting the identification of ground nodes when the connectivity to ground partition is no longer direct.

5 Also, in other embodiments a node 18 is identified as GROUND if its final gradient value has a norm below the predetermined gradient threshold value G , and at least one direct neighbour of it has already been identified as GROUND, and, in addition, the difference between the range of the node (i.e. distance between the sensor and the node) and the range of a direct neighbour of the node
10 that has already been identified as GROUND must be positive (to some tolerance level). In other words, in addition to the criteria used in the above embodiments, in order for a node to be classified as GROUND, a direct neighbour of the node that has already been identified as GROUND must be closer to the sensor. This latter criterion advantageously tends to avoid the label GROUND being propagated from
15 the background of the scene along artificially flat links. Due, for example, to self occlusions a back face of an object is typically not observed in 3D scans. This tends to imply that some of the links in the corresponding mesh "fall off" the apparent side of objects and connect to the background. Such links often correspond to small slopes and as a consequence to flat gradients. Such flat links
20 tend to allow the GROUND label to be propagated from the background onto foreground objects. This tends to materialise as ground points in the middle of object segments. Thus, use of this latter criterion tends to result in a more accurate extraction of the ground.

25 In other embodiments, the flat links that cause ground extraction artefacts, such as that described above, are discarded based on their length. In such embodiments an additional threshold for the length of a flat link is defined.

 Also, in other embodiments the GROUND label is not propagated when the data become too sparse.

30 An optional additional process of analysing and evaluating the segmentation resulting from performing the above described segmentation algorithm may be performed.

- 22 -

Figure 7 is a process flow chart showing certain steps of an example of a method of evaluating the performance of a segmentation algorithm.

In this example, the performance of the segmentation algorithm is evaluated by examining how well the segmentation algorithm segments a particular object (in this example, the tree 10).

At step s12, a segmentation for the tree is manually determined, i.e. a user hand-picking data points that correspond to the tree 10.

In other examples, voxels of data points may be selected instead of single data points, and the evaluation may be performed based on the selected voxels.

Figure 8 is a schematic illustration (not to scale) of a manually determined segmentation of the tree. The boxes of Figure 6 correspond to data points selected by a user.

At step s14, the data points in the segmentation produced by the segmentation algorithm that correspond to those selected by the user (at step s12) are identified.

Thus, a segmentation of the tree 10 as determined by the segmentation algorithm is identified.

Figure 9 is a schematic illustration (not to scale) of a segmentation of the tree determined using the segmentation algorithm.

The data points in Figure 9 (indicated by boxes) are those that correspond to the data points selected by the user. In this example, the data points have been identified as belonging to class-A, class-B, or class-C by the segmentation algorithm. Thus, the data points have been assigned corresponding "identification value" A, B, or C. These identifications are shown in Figure 9. The identification values (A, B, or C) may, for example, indicate separate objects.

In other words, in this example the segmentation algorithm produced a flawed segmentation of the tree 10, and has identified the tree 10 as being made up of multiple partitions. In this example, a partition is a group or cluster of data

- 23 -

points that have the same identification value, that are connected to each other by data points having the same identification value.

At step s16, for each separate partition in the identified segmentation of the tree 10 as determined by the segmentation algorithm, the number of data points in that partition (i.e. partition size) is determined. Also, the identification value of each partition is determined.

In this example, the determined information (i.e. partition size and identification value) may be stored. For example, the information may be stored in a list.

10 The determined information for the segmentation shown in Figure 9 is as follows:

Partition ID value	No. of data points
B	2
A	30
B	12
A	3
C	7

At step s18, the largest partition (i.e. the partition that has the largest number of data points) is identified as a "match".

15 At step s20, the partition(s) not identified as a "match" are identified as "errors".

Thus, in this example the partition that has identification value A, and contains 30 data points, is identified as "match". The remaining four partitions are identified as "errors".

- 24 -

At step s22, a "score" is determined using the number of data points in the partition identified as a "match", and the number of data points in the partitions identified as "errors".

In this example, the score is determined using the following formula:

$$score = \frac{100 * N_{match}}{N_{match} + N_{error}}$$

where: N_{match} is the number of data points in the partition identified as "match"; and

N_{error} is the total number of data points in the partitions identified as "errors".

Thus, the score is the percentage of the total number of data points that are in a partition identified as a match.

In this example, the score is:

$$score = \frac{100 * 30}{54} = 55.6$$

In this example, a score of 95% or more is considered to be a good segmentation of the tree 10. A score of less than 95% but more than 85% is considered to be a mediocre segmentation of the tree. A score of less than 85% is considered to be poor and at about 70% or less is considered to be a failure on behalf of the segmentation algorithm to generate a segmentation of the tree 10. Thus, in this example, the score of 55.6% means that the segmentation algorithm has failed to produce an adequate segmentation for the tree 10. In other examples, different score criteria may be used.

- 25 -

Thus, a method of evaluating the performance of a segmentation algorithm is provided.

An advantage provided by the above described evaluation process is that a value corresponding to how well the segmentation algorithm segments an object is determined. These values may be advantageously used to compare the performance of a plurality of different segmentation algorithms. Also, these values may be advantageously used to compare the performance a segmentation algorithm in multiple operating conditions.

An advantage provided by the above described segmentation evaluation process is that the algorithm can be used as a cost function, e.g. in conjunction with conventional optimisation methods. This may, for example, be used to find optimal parameter values for a segmentation algorithm with free parameters. It may also be advantageously used to determine a most appropriate algorithm, e.g. from a library of possible segmentation methods. For example, a relatively simple approach to optimisation may be used, whereby parameter values for a segmentation algorithm with free parameters are sampled from a parameter space. The segmentation metric is then calculated for that set of sampled values. This may be repeated for different sampled values, and a best performing parameter value set may be selected as optimal. Additionally other more efficient optimisation techniques can be used instead of (uniform) sampling.

In this example, the performance of the segmentation algorithm is evaluated by evaluating how well the segmentation algorithm segments a particular object (in this example, the tree 10). However, in other examples the performance of the segmentation algorithm is evaluated by evaluating how well the segmentation algorithm segments a different object or terrain feature.

In this example, segmentation algorithm evaluation is performed based on data points selected by a user. However, in other examples the data points may be selected by a different appropriate method, for example a different segmentation algorithm.

- 26 -

In other examples, voxels of data points may be selected instead of single data points, and the evaluation may be performed based on the selected voxels. In this case, since all points in a voxel have the same identification value, voxels can be marked as "match" or "error" when the status of any point within the voxel is determined. Voxel match scores tend to be more volatile than the point match score. This is because partitioning of large dense objects like the ground and walls of buildings is more consistent than smaller, sparser objects like trees, people and cars. The quality of partitioning of smaller, sparser, more difficult to partition objects is usually better represented by the voxel score because of less domination by ground and wall points.

It should be noted that certain of the process steps depicted in any of the above described flowcharts may be omitted or such process steps may be performed in differing order to that presented above and shown in the Figures. Furthermore, although all the process steps have, for convenience and ease of understanding, been depicted as discrete temporally-sequential steps, nevertheless some of the process steps may in fact be performed simultaneously or at least overlapping to some extent temporally.

- 27 -

CLAIMS

1. A method for processing an image, the method comprising:
generating a plurality of data points by measuring a value of a parameter at
a plurality of different locations in an area of terrain (4);
5 performing a segmentation process using the plurality of data points to
generate a plurality of image segments, each image segment corresponding to an
object (8, 10) or terrain feature (6);
selecting data points of a particular one of the generated plurality of image
segments; and
10 performing a tracking algorithm using the selected data points to track a
particular object (8, 10) or terrain feature (6) that corresponds to the particular
image segment.
2. A method according to claim 1, wherein the measured parameter is a three
15 dimensional position.
3. A method according to claim 1 or 2, wherein the tracking algorithm
comprises filtering the selected data points.
- 20 4. A method according to any of claims 1 to 3, wherein the tracking algorithm
comprises, for each selected data point, determining a vector that is normal to a
surface of the image segment that corresponds to the particular object or terrain
feature at that data point.
- 25 5. A method according to any of claims 1 to 4, wherein the tracking algorithm
comprises, for each selected data point, determining a range based descriptor for
that data point.

- 28 -

6. A method according to claim 5 when dependent on claim 4, wherein:
the tracking algorithm further comprises, for each selected data point,
determining a coordinate vector; and
a coordinate vector for a selected data point comprises:
5 a value or values corresponding to a time at which that data point
was measured;
a value or values corresponding to the position measured for that
data point;
the range based descriptor determined for that data point;
10 the surface normal determined for that data point; and
an identity value for image segment that that data point belongs to.
7. A method according to any of claims 1 to 6, wherein the tracking algorithm
comprises determining a centroid of the selected data points, and tracking the
15 determined centroid.
8. A method according to any of claims 1 to 7, wherein the segmentation
algorithm comprises:
determining a terrain mesh (12) comprising nodes (18) connected by
20 edges, the nodes (18) corresponding to the data points;
for each node (18), determining a value of a function of the measured
parameter values;
selecting at least one node that corresponds to a determined function value
that is below a threshold value;
25 identifying the selected at least one node as corresponding to the particular
object (8,10) or terrain feature (6); and

- 29 -

identifying a node other than the selected at least one node as corresponding to the particular object (8,10) or terrain feature (6) if:

that node corresponds to a determined function value that is below the threshold value; and

5 that node and the selected at least one node are connected in the terrain mesh (12) by a route, the route comprising either:

a single edge; or

a plurality of edges and one or more nodes, wherein each node on the route corresponds to a determined function value that is below the threshold value.

10

9. A method according to claim 8, wherein the step of identifying a node other than the selected at least one node as corresponding to the particular object (8,10) or terrain feature (6) comprises:

15 identifying the selected at least one node as the current node; and iteratively performing steps (i) to (iii); wherein:

step (i) comprises identifying a node connected in the terrain mesh to the current node by a single edge;

20 step (ii) comprises identifying the node identified at step (i) as corresponding to the particular object (8,10) or terrain feature (6) if that node corresponds to a determined function value that is below the threshold value; and

25 step (iii) comprises identifying the node identified at step (i) as the current node if, at step (ii), the node identified at step (i) is identified as corresponding to the particular object (8,10) or terrain feature (6).

- 30 -

10. A method according to claim 8 or 9, wherein the step of, for each node (18), determining a value of a function using the measured parameter values comprises:

for each edge connected to that node:

5 identifying a node that is at a distance greater than a threshold distance from the node in question in a direction along that edge; and

determining a value of a gradient between the node in question and the identified node; and

determining a value of a function of the gradient values.

10

11. A method according to claim 10, wherein the function of the gradient values is a maximum of the absolute values of the gradient values.

12. A method according to any of claims 8 to 11, wherein the method further
15 comprises identifying as corresponding to a different object (8, 10) or terrain feature (6) to the particular object (8, 10) or terrain feature (6), each of the nodes in the terrain mesh (12) not identified as corresponding to the particular object (8, 10) or terrain feature (6).

20 13. Apparatus for processing an image, the apparatus comprising:

a sensor (2) arranged to measure a value of a parameter at a plurality of different locations in an area of terrain (4) so as to generate a plurality of data points; and

one or more processors (3) arranged to:

25 perform a segmentation process using the plurality of data points to generate a plurality of image segments, each image segment corresponding to an object (8, 10) or terrain feature (6); and

- 31 -

perform a tracking algorithm using data points of one of the generated plurality of image segments so as to track a particular object (8, 10) or terrain feature (6) that corresponds to that image segment.

5 14. 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 12.

10 15. A machine readable storage medium storing a computer program or at least one of the plurality of computer programs according to claim 14.

1/9




FIG. 1

2/9




FIG. 2




FIG 3




FIG 4

5/9






FIG. 5



7/9



8/9




FIG. 8

9/9




FIG. 9

INTERNATIONAL SEARCH REPORT

International application No.

PCT/AU2012/000248

A. CLASSIFICATION OF SUBJECT MATTER

Int. Cl.

G06T 17/05 (2011.01)

G06K 9/46 (2006.01)

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, surface, track, detect, identify, segment, section, area, region, pixel, three dimensional, position and like words.

Google Patents, Google Scholar & Patent-Lens websites: 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.
X Y	US 6072889 A (DEAETT et al.) 6 June 2000 See abstract; column 2, lines 46-64; column 8, lines 59-62; and figure 17.	1, 3, 7 and 13-15 2 and 4-6
Y	MOOSMANN, F. et al., 'Motion Estimation from Range Images in Dynamic Outdoor Scenes', 2010 IEEE International Conference on Robotics and Automation, May 3-8, 2010, pages 142-147.	2 and 4-6



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

19 April 2012

Date of mailing of the international search report

24 April 2012

Name and mailing address of the ISA/AU

AUSTRALIAN PATENT OFFICE
PO BOX 200, WODEN ACT 2606, AUSTRALIA
E-mail address: pct@ipaustralia.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/AU2012/000248

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	6072889	AU	19963/99
		EP	1036340
		JP	2001525548
		WO	9928759
Due to data integration issues this family listing may not include 10 digit Australian applications filed since May 2001.			
END OF ANNEX			