



- (51) **International Patent Classification:**
H04N 13/02 (2006.01) **G03B 35/00** (2006.01)
H04N 5/225 (2006.01)
- (21) **International Application Number:**
PCT/US2012/020681
- (22) **International Filing Date:**
9 January 2012 (09.01.2012)
- (25) **Filing Language:** English
- (26) **Publication Language:** English
- (30) **Priority Data:**
13/017,474 31 January 2011 (31.01.2011) US
- (71) **Applicant (for all designated States except US):** **MICROSOFT CORPORATION** [US/US]; One Microsoft Way, Redmond, Washington 98052-6399 (US).
- (72) **Inventors:** **NEWCOMBE, Richard**; c/o Microsoft Corporation, LCA - International Patents, One Microsoft Way, Redmond, Washington 98052-6399 (US). **IZADI, Shahram**; c/o Microsoft Corporation, LCA - International

Patents, One Microsoft Way, Redmond, Washington 98052-6399 (US). **MOLYNEAUX, David**; c/o Microsoft Corporation, LCA - International Patents, One Microsoft Way, Redmond, Washington 98052-6399 (US). **HILLIGES, Otmar**; c/o Microsoft Corporation, LCA - International Patents, One Microsoft Way, Redmond, Washington 98052-6399 (US). **KIM, David**; c/o Microsoft Corporation, LCA - International Patents, One Microsoft Way, Redmond, Washington 98052-6399 (US). **SHOTTON, Jamie Daniel Joseph**; c/o Microsoft Corporation, LCA - International Patents, One Microsoft Way, Redmond, Washington 98052-6399 (US). **KOHLI, Pushmeet**; c/o Microsoft Corporation, LCA - International Patents, One Microsoft Way, Redmond, Washington 98052-6399 (US). **FITZGIBBON, Andrew**; c/o Microsoft Corporation, LCA - International Patents, One Microsoft Way, Redmond, Washington 98052-6399 (US). **HODGES, Stephen Edward**; c/o Microsoft Corporation, LCA - International Patents, One Microsoft Way, Redmond, Washington 98052-6399 (US). **BUTLER, David Alexander**; c/o Microsoft Corporation, LCA - International Patents, One Microsoft Way, Redmond, Washington 98052-6399 (US).

[Continued on next page]

(54) **Title:** MOBILE CAMERA LOCALIZATION USING DEPTH MAPS




Fig. 3

(57) **Abstract:** Mobile camera localization using depth maps is described for robotics, immersive gaming, augmented reality and other applications. In an embodiment a mobile depth camera is tracked in an environment at the same time as a 3D model of the environment is formed using the sensed depth data. In an embodiment, when camera tracking fails, this is detected and the camera is relocalized either by using previously gathered keyframes or in other ways. In an embodiment, loop closures are detected in which the mobile camera revisits a location, by comparing features of a current depth map with the 3D model in real time. In embodiments the detected loop closures are used to improve the consistency and accuracy of the 3D model of the environment.



(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).

Declarations under Rule 4.17:

- as to applicant's entitlement to apply for and be granted a patent (Rule 4.17(ii))
- as to the applicant's entitlement to claim the priority of the earlier application (Rule 4.17(iii))

Published:

- without international search report and to be republished upon receipt of that report (Rule 48.2(g))

MOBILE CAMERA LOCALIZATION USING DEPTH MAPS

BACKGROUND

- [0001] Mobile camera localization involves finding the position and orientation of a camera moving in its environment and is useful for many applications such as robotics, immersive gaming, augmented reality, architecture, planning, robotics, engineering prototyping, vehicle navigation, medical applications and other problem domains. Existing approaches are limited in accuracy, robustness and speed. For many applications accurate camera localization is required in real time, for example, in order that a robot may successfully move about in its environment.
- [0002] Many previous approaches to mobile camera localization have used color video cameras rather than depth cameras. Typically color video cameras give high resolution and accuracy and the rich color information allows visual features to be detected in the video images. Information available from depth cameras may be noisy and inaccurate depending on the type of environment and the type of depth camera used. Depth cameras capture images referred to herein as depth maps where each pixel relates to an absolute or relative distance from the depth camera to a point in the camera's environment. It may be harder to detect features in depth maps as compared with color images due to the difference in the information available.
- [0003] Some previous approaches to mobile camera localization have involved creating a map of a mobile camera's environment at the same time as tracking the camera's position and orientation with respect to that map. This is known as simultaneous localization and mapping (SLAM).
- [0004] The embodiments described below are not limited to implementations which solve any or all of the disadvantages of known mobile camera localization systems.

SUMMARY

[0005] The following presents a simplified summary of the disclosure in order to provide a basic understanding to the reader. This summary is not an extensive overview of the disclosure and it does not identify key/critical elements of the invention or delineate the scope of the invention. Its sole purpose is to present a selection of concepts disclosed herein in a simplified form as a prelude to the more detailed description that is presented later.

[0006] Mobile camera localization using depth maps is described for robotics, immersive gaming, augmented reality and other applications. In an embodiment a mobile depth camera is tracked in an environment at the same time as a 3D model of the environment is formed using the sensed depth data. In an embodiment, when camera tracking fails, this is detected and the camera is relocalized either by using previously gathered keyframes or in other ways. In an embodiment, loop closures are detected in which the mobile camera revisits a location, by comparing features of a current depth map with the 3D model in real time. In embodiments the detected loop closures are used to improve the consistency and accuracy of the 3D model of the environment.

[0007] Many of the attendant features will be more readily appreciated as the same becomes better understood by reference to the following detailed description considered in connection with the accompanying drawings.

DESCRIPTION OF THE DRAWINGS

[0008] The present description will be better understood from the following detailed description read in light of the accompanying drawings, wherein:

FIG. 1 is a schematic diagram of a person in a room holding a mobile depth camera which may be used for real-time camera tracking and optionally also to produce a dense 3D model or map of the room;

FIG. 2 is a plan view of a floor of a building which is being explored by a person holding a mobile depth camera;

FIG. 3 is a schematic diagram of a mobile depth camera connected to a real-time camera tracking system, a dense 3D model formation system and a game system;

- FIG. 4 is a flow diagram of a method at a real-time tracker;
FIG. 5 is a schematic diagram of a relocalization engine;
FIG. 6 is a flow diagram of a method of relocalization without keyframes;
FIG. 7 is a flow diagram of a method of relocalization with keyframes;
5 FIG. 8 is a flow diagram of a method of loop closure;
FIG. 9 is a schematic diagram of an example frame alignment engine;
FIG. 10 is a flow diagram of an iterative process for camera tracking;
FIG. 11 is a flow diagram of more detail of part of the iterative process of
FIG. 5 for computing pairs of corresponding points;
10 FIG. 12 is a flow diagram of a process for computing pairs of corresponding
points using predictions from a dense 3D model;
FIG. 13 is a flow diagram of a process for computing and minimizing a
point-to-plane error metric for use in the iterative process of FIG. 10;
FIG. 14 illustrates an exemplary computing-based device in which
15 embodiments of a camera localization system may be implemented.

Like reference numerals are used to designate like parts in the accompanying drawings.

DETAILED DESCRIPTION

[0009] The detailed description provided below in connection with the
20 appended drawings is intended as a description of the present examples and is not
intended to represent the only forms in which the present example may be
constructed or utilized. The description sets forth the functions of the example and
the sequence of steps for constructing and operating the example. However, the
same or equivalent functions and sequences may be accomplished by different
25 examples.

[0010] Although the present examples are described and illustrated herein as
being implemented in a real-time camera tracking system using depth images
obtained from a mobile depth camera which emits and captures infra-red light, the
system described is provided as an example and not a limitation. As those skilled
30 in the art will appreciate, the present examples are suitable for application in a
variety of different types of real-time camera tracking systems including but not

limited to those using depth information obtained from stereo cameras and those using depth information obtained by emitting and capturing other types of electromagnetic radiation.

[0011] The term "image element" is used in this document to refer to a pixel,
5 group of pixels, voxel or other higher level component of an image.

[0012] The term "dense 3D model" is used in this document to refer to a representation of a three dimensional scene comprising objects and surfaces where that representation comprises detail about image elements of that scene. In contrast sparse 3D models may comprise frame-based representations of objects. A dense
10 3D model may be transformed to a sparse 3D model such as a polygon mesh representation, or other representation in a manner which reduces redundancy and memory required to store the 3D model. A example dense 3D model may be one in which all or many points from an incoming depth map are used to describe surfaces in the environment. A sparse model would take a subset of the points to speed up
15 calculations and reduce memory footprint.

[0013] FIG. 1 is a schematic diagram of a person 100 standing in a room and holding a mobile depth camera 102 which in this example also incorporates a projector which is projecting the image of a cat 108 into the room. The room contains various objects 106 such as a chair, door, window, plant, light and another
20 person 104. Many of the objects 106 are static although some of the objects such as person 104 may move. As the person moves around the room the mobile depth camera captures images which are used by a real-time camera tracking system 112 to monitor the location and orientation of the camera in the room. The real-time camera tracking system 112 may be integral with the mobile depth camera 102 or
25 may be at another location provided that it is able to receive communication from the mobile depth camera 102, either directly or indirectly. For example, the real-time camera tracking system 112 may be provided at a personal computer, dedicated computer game apparatus, or other computing device in the room and in wireless communication with the mobile depth camera 102. In other examples the
30 real-time camera tracking system 112 may be elsewhere in the building or at another remote location in communication with the mobile depth camera 102 using

a communications network of any suitable type. The mobile depth camera 102 is also in communication with a dense 3D model 110 of the environment (which in this case is a 3D model of the room) or another type of map of the environment. For example, images captured by the mobile depth camera 102 are used to form and
5 build up the dense 3D model of the environment as the person moves about the room. The real-time camera tracking system 112 may track the position of the camera in relation to the 3D model or map of the environment 110. The outputs of the real-time camera tracking system 112 and dense 3D model or map 110 may be used by a game system or other application although that is not essential. For
10 example, a projector at the mobile depth camera 102 may be arranged to project images depending on the output of the real-time camera tracking system 112 and 3D model 110.

[0014] FIG. 2 is a plan view of a floor 200 of a building. A person 202 holding a mobile depth camera 204 is moving around the floor as indicated by
15 dotted arrows 208. The person walks along a corridor 206 past rooms and furniture 210. The real-time camera tracking system 112 is able to track the position of the mobile depth camera 204 as it moves and a 3D model or map of the floor is formed. It is not essential for a person 202 to carry the mobile depth camera 204. In other examples the mobile depth camera 204 is mounted on a robot or vehicle. This also
20 applies to the example of FIG.1.

[0015] FIG. 3 is a schematic diagram of a mobile environment sensor 300 for use with a real-time camera tracker 316, dense model formation system 324 and optionally a game system 332. The mobile environment sensor 300 comprises a depth camera 302 which is arranged to capture sequences of depth images of a
25 scene. Each depth image or depth map frame 314 comprises a two dimensional image in which each image element comprises a depth value such as a length or distance from the camera to an object in the captured scene which gave rise to that image element. This depth value may be an absolute value provided in specified units of measurement such as meters, or centimeters or may be a relative depth
30 value. In each captured depth image there may be around 300,000 or more image elements each having a depth value. The frame rate is high enough to enable the

depth images to be used for working robotics, computer game or other applications. For example, at least 20 frames per second.

[0016] The depth information may be obtained using any suitable technique including but not limited to, time of flight, structured light, stereo images. In some
5 examples the depth camera is able to organize the depth information into Z layers that are perpendicular to a Z axis extending along a line of sight of the depth camera.

[0017] The mobile environment sensor 300 may also comprise an emitter 304 arranged to illuminate the scene in such a manner that depth information may
10 be ascertained by the depth camera 302. For example, in the case that the depth camera 302 is an infra-red (IR) time-of-flight camera, the emitter 304 emits IR light onto the scene, and the depth camera 302 is arranged to detect backscattered light from the surface of one or more objects in the scene. In some examples, pulsed infrared light may be emitted from the emitter 304 such that the time between an
15 outgoing light pulse and a corresponding incoming light pulse may be detected by the depth camera and measure and used to determine a physical distance from the environment sensor 300 to a location on objects in the scene. Additionally, in some examples, the phase of the outgoing light wave from the emitter 304 may be compared to the phase of the incoming light wave at the depth camera 302 to
20 determine a phase shift. The phase shift may then be used to determine a physical distance from the mobile environment sensor 300 to a location on the objects by analyzing the intensity of the reflected beam of light over time via various techniques including, for example, shuttered light pulse imaging.

[0018] In another example, the mobile environment sensor 300 can use
25 structured light to capture depth information. In such a technique patterned light (e.g. light displayed as a known pattern such as a grid or stripe pattern) may be projected onto a scene using the emitter 304. Upon striking the surface of objects in the scene the pattern becomes deformed. Such a deformation of the pattern is captured by the depth camera 302 and analyzed to determine an absolute or relative
30 distance from the depth camera 302 to the objects in the scene.

[0019] In another example, the depth camera 302 comprises a pair of stereo camera such that visual stereo data is obtained and resolved to generate relative depth information. In this case the emitter 304 may be used to illuminate the scene or may be omitted.

5 [0020] In some examples, in addition to the depth camera 302, the mobile environment sensor 300 comprises a color video camera referred to as an RGB camera 306. The RGB camera 306 is arranged to capture sequences of images of the scene at visible light frequencies.

[0021] The mobile environment sensor 300 may comprise an orientation
10 sensor 308 such as an inertial measurement unit (IMU), accelerometer, gyroscope, compass or other orientation sensor 308. However, it is not essential to use an orientation sensor. The mobile environment sensor 300 may comprise a location tracking device such as a GPS although this is not essential.

[0022] The mobile environment sensor may comprise a projector 312 as
15 mentioned above with reference to FIG. 1 although this is not essential.

[0023] The mobile environment sensor also comprises one or more processors, a memory and a communications infrastructure as described in more detail below.

[0024] The mobile environment sensor may be provided in a housing which
20 is shaped and sized to be hand held by a user or worn by a user. In other examples the mobile environment sensor is sized and shaped to be incorporated or mounted on a vehicle, toy or other movable apparatus.

[0025] The mobile environment sensor 300 is connected to a real-time tracker 316. This connection may be a physical wired connection or may use
25 wireless communications. In some examples the mobile environment sensor 300 is connected indirectly to the real-time tracker over one or more communications networks such as the internet.

[0026] The real-time tracker 316 is computer implemented using a general purpose microprocessor controlling one or more parallel computing units such as
30 graphics processing units (GPUs), vector machines, multi-core processors or other parallel computing devices. It comprises a frame alignment engine 318 and

optionally a loop closure engine 320 and a relocalization engine 322. The real-time tracker 316 takes depth map frames 314 from the depth camera 302, and optionally also input from the mobile environment sensor 300, optional map data 334 and optional data from a game system 332. The real-time tracker operates to place the
5 depth map frames into alignment in order to produce a real-time series 328 of six degree of freedom pose estimates of the depth camera 302. It may also produce transformation parameters (also referred to as registration parameters) for transforms between pairs of depth map frames. In some examples the real-time tracker operates on pairs of depth map frames 314 from the depth camera. In other
10 examples, the real-time tracker 216 takes a single depth map frame 314 and aligns that with a dense 3D model 326 of the scene rather than with another depth map frame 314.

[0027] For example, in some embodiments the real-time tracker 316 provides output to a dense 3D model formation system 324 which uses that
15 information together with the depth map frames 314 to form and store a dense 3D model of the scene or environment in which the mobile environment sensor 300 is moving. For example, in the case of FIG. 1 the 3D model would be a 3D model of the surfaces and objects in the room. In the case of FIG. 2 the 3D model would be a 3D model of the floor of the building. The dense 3D model 326 may be stored in
20 GPU memory or in other ways.

[0028] The mobile environment sensor 300 may be used in conjunction with a game system 332 which is connected to a display 330. For example, the game may be a golf game, boxing game, motor car racing game or other type of computer game. Data from the game system 332 such as the game state or meta data about
25 the game may be provided to the real-time tracker 316. Also, information from the real-time tracker may be used by the game system 332 to influence how the game proceeds. Information from the 3D model may also be used by the game system 332 to influence how a game proceeds.

[0029] Map data 334 is optionally available to the real-time tracker 316. For example, this may be an architect's drawing of the environment (e.g. room or floor of building), the location of landmarks known in the environment, a map of the environment available from another source.

5 [0030] The frame alignment engine 318 of the real-time tracker is arranged to align pairs of depth map frames or a depth map frame and an estimate of a depth map frame from the dense 3D model. It uses an iterative process which is implemented using one or more graphics processing units in order that the frame alignment engine operates in real-time. More detail about the frame alignment
10 engine is given below with reference to FIG. 9.

[0031] The processing performed by the real-time tracker 316 and/or the dense 3D model formation system 324 can, in one example, be executed remotely from the location of the mobile environment capture device 300. For example, the mobile environment capture device 300 can be connected to (or comprise) a
15 computing device having relatively low processing power, and which streams the depth images over a communications network to a server. The server has relatively high processing power, and performs the computationally complex tasks of the real-time tracker 316 and/or the dense 3D model formation system 324. The server can return a rendered image of the dense reconstruction per-frame to provide an
20 interactive experience to the user, and also return the final dense 3D reconstruction on completion of the model, for subsequent local use (e.g. in a game). Such an arrangement avoids the need for the user to possess a high-powered local computing device.

[0032] The loop closure engine is arranged to detect when the mobile
25 environment sensor has moved in a loop so that the scene depicted in the current depth frame is at least partially overlapping with that of a previous depth frame which is not the immediately preceding depth frame. When a loop is closed the mobile environment sensor revisits somewhere it has been before. For example, this may occur when a user walks around the whole floor of the building in FIG. 2
30 and reaches the starting point again. It may also occur when a user moves around a room behind some furniture and out again to the original start position, or close to

that original start position. Thus the size of the loop will vary depending on the environment that the mobile environment sensor is in and on the application domain. For example, in the case of a person walking around the floor of a building (as in FIG. 2) the loop may be tens of meters long. In the case of a person
5 moving a camera around a single room (as in FIG. 1) the loop may be less than 10 meters long. In the case of a robot on a toy motor car driving around a child's toy racing track the loop may be less than 1 meter long. Detecting when a loop closure occurs is useful because it enables accumulated errors in the 3D model to be identified. For example, the mobile environment sensor may have captured depth
10 and/or color video images identified as being at the start and the end of a loop but the camera position and orientation calculated for each of those images may not be consistent. Once these errors are identified they may be reduced so improving the accuracy and consistency of the 3D model or map of the environment. Also, gross errors or disorientation of the camera tracking may be recovered from. Loop
15 closure errors can be thought of as comprising both localized errors and global errors. Global loop closure errors (also referred to as drift) comprise errors in the combined 6 degree of freedom pose estimates of the camera over time. Localized errors may occur in 6 degree of freedom pose estimates of the camera from individual frames, or sequences of largely overlapping frames. , Accurately
20 detecting when a loop closure occurs in real-time is not straightforward. Once a loop closure has been detected, any loop closure error is to be identified and reduced and this is difficult to achieve in real time whilst achieving a high quality result. In addition, global loop closure errors may be counteracted or treated in a different manner from local loop closure errors.

25 [0033] The relocalization engine 322 is arranged to deal with the situation where the real-time tracker loses the current location of the mobile environment sensor 300 and relocalizes or finds the current location again. The real-time tracker may lose the current location of the mobile environment sensor 300 for many reasons. For example, rapid camera motion, occlusion and motion blur. In the case
30 that depth maps are used for real-time tracking, tracking failures may also occur where the environment has little fine detail and comprises mainly planar surfaces.

In the case where depth cameras which use a structured light emitter tracking failures may occur where the environment comprises highly reflective surfaces such as computer screens and other reflective surfaces which produce poor depth information. If tracking failures are not detected, erroneous tracking information may potentially be used to increment the 3D model or map of the environment. This may lead to corruption of the 3D model or map. Fast and accurate detection of tracking failures is beneficial because it can be used to prevent corruption of the 3D model as well as to trigger a tracking recovery process.

[0034] FIG. 4 is a flow diagram of a process at the real-time tracker 316 of

FIG. 3. Real-time camera tracking is established 400 using the frame alignment engine 318 as described with reference to FIG.s. 9 to 13 below. Depth map frames 314 and the output of the frame alignment engine (registration parameters and camera orientation and position) are used 402 to refine a dense 3D model 326 of the environment in which the mobile environment sensor 300 is moving. If a tracking failure is detected 404 then relocalization 406 of the mobile environment sensor 300 takes place and the process continues at step 402 with incrementing the dense 3D model as more depth map frames are captured. In this way, ongoing camera tracking and ongoing refinement of the 3D model is paused once a tracking failure is detected and until relocalization is successful. If a loop closure error is detected 408 then the dense 3D model itself is modified to take into account that error. After addressing the loop closure error the process continues at step 402 with incrementing the dense 3D model. In this way the dense 3D model of the environment is preserved when tracking failures occur and its accuracy and consistency is improved as a result of detecting loop closures.

[0035] FIG. 5 is a schematic diagram giving more detail of the relocalization engine 322 of FIG. 3. It comprises a tracking failure detector 500 which operates in real time. The tracking failure detector 500 may comprise thresholds 502 which are pre-configured depending on the application domain concerned. For example, a tracking failure may occur when the registration parameters output by the frame alignment engine change by more than a threshold amount. In another example, a tracking failure may occur when the SE3 matrix output by the frame alignment

engine changes by more than a threshold amount between successive outputs. A tracking failure may be detected when a failure in convergence of an iterative process used to track the position and orientation of the mobile depth camera occurs. The tracking failure detector 500 may comprise a motion heuristic 504
5 stored as one or more rules, criteria, thresholds or other conditions. For example, the motion heuristic may be a rule that the mobile environment sensor 300 is moving at a constant velocity with white noise acceleration. Any output from the frame alignment engine which is not consistent with this motion heuristic may trigger a tracking failure. In another example, the motion heuristic may be that the
10 mobile environment sensor 300 moves with a random walk with a maximum linear velocity set at walking speed of an average human (in the case that the mobile environment sensor is worn by a user) and with an arbitrarily high rotational velocity. Any output from the frame alignment engine which is not consistent with this motion heuristic may trigger a tracking failure. Any combination of one or
15 more of these ways of detecting a tracking failure may be used. In an example, a combination of three aspects of violation is used to detect a tracking failure according to violation of a motion heuristic, violation of rotation thresholds and/or when failure of an iterative process used to track the position and orientation of the mobile depth camera occurs due to failure to converge or converging with too large
20 a residual error.

[0036] The tracking failure detector 500 may comprise a component which is arranged to detect a tracking failure when the frame alignment engine detects a degeneracy 506. For example, as described below, the frame alignment engine finds pairs of corresponding points between pairs of depth maps by using an
25 iterative process that comprises optimization of an error metric. During optimization of that error metric a degeneracy may be detected when that optimization process fails to give any solution. For example, this failure may be because of image elements in the depth maps which do not have depth values because of reflective surfaces in the environment, because the environment is
30 highly planar, if spherical or cylindrical surfaces are present, or for other reasons.

[0037] The relocalization engine also comprises a relocalization process 508 which may use keyframes or may operate without keyframes. Keyframes are depth maps or color video frames gathered previously by the real time tracker or artificially generated. Tracking is recovered by finding a key frame that is a good match to the current camera view and this provides an initial pose estimate from which tracking can recommence. In some cases the keyframes are marginalized into an SE3 matrix and this reduces memory requirements. The keyframe points may be recovered from the dense 3D model by ray casting when required.

[0038] An example in which the relocalization process 508 operates without keyframes is now given with reference to FIG. 6. A track distribution of camera poses is computed 600 for example, using a motion model for the mobile environment sensor. This motion model may predict that the mobile environment sensor is within a steadily expanding sphere centered about its last known position and with any orientation. A sample camera pose is selected 602 from the track distribution at random or in any other manner. The frame alignment engine is then used 604 with the sampled camera pose to align the current depth map and previous depth map to obtain registration parameters as described in more detail below. If convergence is reached 606, that is, if the registration parameters are consistent with the last known position, then the camera has been successfully relocalized 608. Otherwise, another sample is taken 602 and the process repeats.

[0039] Another example in which the relocalization process 508 operates without keyframes is now given.

[0040] A fast clustering algorithm such as a random decision forest is applied to patches of the current depth map and to patches of a plurality of previous depth maps obtained from the 3D model of the environment. The previous depth maps may be obtained from the 3D model of the environment by using a ray casting technique to render depth maps from the 3D model or in any other way. Each leaf node in the random decision forest represents a texton. Textons are texture features of a color image but in this example, depth maps are used and so the textons are texture features of a depth map. A plurality of patches of the current depth map is selected at random or in any other manner. Each patch is a plurality

of adjacent image elements. Patches are also selected from each of the plurality of previous depth maps. Each patch is processed to a leaf of the random decision forest to identify a texton for that patch. A histogram is then built of all the textons found in an image.

5 [0041] The relocalization process selects a previous depth map which is similar to the current depth map in terms of a histogram of the textons output by the random decision forest classifier. The camera pose associated with the selected depth map is then used as the current camera poses and the camera is relocalized. The random decision forest clustering algorithm is trained offline or using a
10 background process during use of the real time tracker.

[0042] FIG. 7 is a flow diagram of a method of using keyframes for relocalization. Keyframes are gathered 702 during real time tracking and stored in association with the 3D model. For example, the keyframes are color video frames captured by an RGB camera 306 at the mobile environment sensor 300. In another
15 example, the keyframes are depth map frames 314 captured by the depth camera 302. Each keyframe has an associated camera pose which is calculated by the frame alignment engine 318 during real time tracking. Each keyframe is associated with a location in the 3D model 326 according to the camera pose. The keyframes are stored in association with the 3D model, for example, by including a reference
20 in the 3D model at the location specified by the camera pose for the keyframe. In some examples, the 3D model is stored in a cube of GPU memory and the keyframes are hung in this cube of memory according to the locations specified by the camera poses for the keyframes. In some cases the keyframes are stored as their associated SE3 matrix and this reduces memory requirements. The keyframe depth
25 points may be recovered from the dense 3D model by ray casting when required.

[0043] The number of keyframes that are gathered depends on the memory and processing capacity available as well as on the application domain. However, tens of thousands of keyframes may be stored. Selection of frames to keep as keyframes is achieved using random selection, by selecting every nth frame, by
30 selecting frames which comprise previously unseen visual features, using combinations of these approaches, or in any other way. In an example, the system

waits a minimum number of frames between keyframes and records a new keyframe if that frame corresponds to a camera pose which is at least a specified amount away from any of the existing keyframes in terms of translation and rotation.

5 [0044] When a tracking failure is detected 702 the relocalization engine selects 704 one of the keyframes which is similar to the current camera view (depth map or RGB image). The measure of similarity between the keyframe and the current camera view may be of any suitable type. For example, features such as edges, corners, blobs, lines may be compared. Other measures of similarity may
10 also be used such as comparing results of an object segmentation process applied to the keyframe and the current camera view. In the case that the keyframes are depth maps, a human body pose detector may be applied to the keyframe and the current view and the results compared. A geometry estimation and/or semantic image labeling process (such as a machine learning classifier) may be applied to the
15 current and previous frames. It is then possible to relocalize by selecting frames which have similar geometry (arrangement and locations of vertical and horizontal surfaces) and/or semantic labeling (arrangement of objects) compared to the current frame.

[0045] Selecting a key frame is a complex and time consuming task as there
20 are so many keyframes to search through. In some examples, the search is made of keyframes which are located around the last known camera position and then the search space may be gradually widened until a good match is found.

[0046] Once a keyframe has been selected it is used 706 to provide an initial pose estimate from which tracking can recommence.

25 [0047] A method at the loop closure engine 320 of FIG. 3 is now described with reference to FIG. 8. A loop closure detection process 804 takes input from the current depth map 800 and the 3D model 802 of the environment. In some cases the loop closure detection process 804 takes input from a current color image captured by the RGB camera 306 although this is not essential. The loop closure
30 detection process 804 uses any suitable method of loop closure detection such as a feature matching process 806, a bag of words process 808 or other loop closure

detection process. As mentioned above, loop closure occurs when the mobile environment sensor revisits a location. Bag of words processes are able to efficiently recognize when two images show the same scene and so detect loop closure. A bag of words process generates a dictionary (either dynamically or
5 offline) of feature descriptors (e.g. lines, edges, corners, blobs, SIFT features, SURF features) and uses the dictionary to label features found in images. The images can then be compared very quickly to see if they have many features in common and hence may be of the same object or location. A feature matching process is any process which identifies features that are the same in two images
10 without using a dictionary.

[0048] Once loop closure is detected any loop closure error is found and may be accounted for. For example, the loop closure identifies that part of the 3D model and the current depth map have the same camera location. However because of tracking errors this may not be the case. In addition, the type of the loop closure
15 error may be identified. For example, whether the loop closure is a global one such as being the result of drift in the combined 6 degree of freedom pose estimated of the camera, or whether the error is a local one which is the result of poor 6 degree of freedom pose estimates from an individual frame or sequence of largely overlapping frames. The loop closure error may be a combination of global and
20 local loop closure errors. A loop closure procedure 812 is followed to reduce this error and to improve the consistency and accuracy of the 3D model. Any suitable loop closure procedure may be used such as optimizing an energy function.

[0049] FIG. 9 is a schematic diagram of the frame alignment engine 318 of FIG. 3 in more detail. The frame alignment engine 908 is computer implemented at
25 a computing device having at least a CPU and one or more GPUs. It comprises a plane extraction component 910 which is optional and an iterative closest point process 912. The iterative closest point process uses projective data association and a point-to-plane error metric as described in more detail below. The frame alignment engine receives the current depth map 900 from the depth camera. This
30 is also referred to as the destination depth map. In some examples it also receives a source depth map 902 which is the previous depth map frame from the depth

camera. In other examples, the frame alignment engine takes a dense surface model estimate 906 of the source depth map. The output of the frame alignment engine is a set of registration parameters of a transform for aligning the current and source frames (or frame estimate). In some examples these registration parameters are provided as a six degree of freedom (6DOF) pose estimate in the form of an SE_3 matrix describing the rotation and translation of the depth camera 302 relative to real-world coordinates. More formally, this transformation matrix can be expressed as:

$$\mathbf{T}_k = \begin{bmatrix} \mathbf{R}_k & \mathbf{t}_k \\ 0^\top & 1 \end{bmatrix} \in \mathbf{SE}_3$$

10 [0050] Where $\mathbf{T}_k$ is the transformation matrix for depth image frame k , $\mathbf{R}_k$ is the camera rotation for frame k , $\mathbf{t}_k$ is the camera translation at frame k , and Euclidean group $\mathbf{SE}_3 := \{\mathbf{R}, \mathbf{t} \mid \mathbf{R} \in SO_3, \mathbf{t} \in \mathbb{R}^3\}$. Coordinates in the camera space (i.e. from the camera perspective) can be mapped to real-world coordinates by multiplying by this transformation matrix. . However, the registration parameters may be provided in any suitable form. These registration parameters are used by the real time tracker 316 to produce the real-time series of 6 degree of freedom pose estimates of the depth camera.

[0051] FIG. 10 is a flow diagram of an example iterative process at a frame alignment engine. An initial estimate of the registration parameters is formed 1000. 20 These are the registration parameters of a transform for aligning the current and source frames. This initial estimate is formed in any suitable manner. For example, one or more of the following sources of information may be used to form the initial estimate: game state, game metadata, map data, RGB camera output, orientation sensor output, GPS data. In another example, the initial estimate is formed by 25 predicting where the camera is using information about the previous motion path of the camera. For example, the camera may be assumed to have a constant velocity or a constant acceleration. The motion path of the camera from time 0 to time $t-1$ may be used to estimate where the camera will be at time t and thus obtain an estimate of the registration parameters.

[0052] Using the initial estimate, pairs of corresponding points between the current and source frames (depth maps or estimated depth maps) are computed 1002. A pair of corresponding points is a point from one depth map and a point from another depth map, where those points are estimated to have arisen from the same real world point in a scene. The term "point" is used here to refer to a pixel, or a group or patch of neighboring pixels. This correspondence problem is very difficult because of the huge number of possible combinations of points. Previous approaches using color or grey-scale images have addressed this problem by identifying shapes such as lines, edges, corners or the like in each image and then trying to match those shapes between the pair of images. In contrast, the embodiments described herein identify corresponding points without the need to find shapes in the depth maps. More detail about how the corresponding points are computed is given below with reference to FIG. 11. An updated estimate of the registration parameters is calculated 1004 which optimizes an error metric applied to the computed corresponding points.

[0053] A check is made to assess whether convergence has been reached 1006. If so, there is little or no change in the updated estimate and the registration parameters are output 1008. If not, the iterative process repeats as indicated in FIG. 10.

[0054] With reference to FIG. 11 more detail is now given about how the pairs of corresponding points are calculated. In some embodiments sample points are taken 1100 from either or both of the current and source depth maps and those sample points are used as candidates to find pairs of corresponding points from. Sampling may be achieved by randomly selecting a specified proportion of the points. In another embodiment sampling is achieved in a manner which takes into account surface normals of the points. For example, a surface normal is calculated for each point (as described in more detail below) and a histogram created with a plurality of bins for different ranges of surface normal values. Sampling is carried out so that a uniform sampling across the bins is achieved.

30

[0055] By using sampling the benefit that computational costs are reduced is achieved. However, a risk is that the accuracy and robustness of the process is reduced when sampling is used. This is because the sample points may not provide a good indication of the depth map that the samples were taken from. For example, the samples may enable the process to find a set of corresponding points which the process identifies as a solution but which in fact represent a locally optimal solution rather than a globally optimal solution.

[0056] Having said that, it is not essential to use any sampling. The process is also workable and gives good results when all the available points are used. In this case the GPU implementation described herein allows the process to operate in real-time for all the points in each depth map which may be as many as 300,000 or more. In the example described below with reference to FIG. 11 the process is described as using sampling. However, the process of FIG. 11 is also applicable where no sampling is carried out.

[0057] As indicated in FIG. 11, surface normals of the sample points (or each available point in the case that no sampling is done) are computed 1102. For example, this is achieved for a given point by finding the four (or more) nearest neighbor points in the depth map and computing a surface patch which incorporates those neighbors and the point itself. A normal to that surface patch is then calculated at the location of the point.

[0058] A process of finding 1104 corresponding pairs of points is then followed. This is now described in the case that the source and current depth maps are available without the use of a dense 3D model. For each sampled source point from the source depth map, a ray is projected 1106 from the camera location associated with the source depth map, through the sampled source point and onto a destination point in the destination depth map. In some cases the destination point may be in front of the sampled source point along the projected ray. This projection process may be referred to as "projective data association". A search 1108 is then made for candidate corresponding points around and including the destination point. For example, the search is for points which have surface normals that are compatible with the surface normal of the sampled source point. Surface

normals are said to be compatible if they are within a specified range of one another and which are within a specified Euclidean distance of the destination point. For example, this specified range is user configurable.

[0059] One or more candidate corresponding points are found as a result of this search. From those candidate corresponding points a single point is selected 1110 to form a pair with the source point. This selection is made on the basis of a distance metric. For example, a Euclidean distance is calculated between the source point and each of the candidate corresponding points. The pair which gives the smallest Euclidean distance is then selected. The process of box 1104 is then repeated for each of the sampled source points or in the case that no sampling is done, for each of the available points of the source depth map.

[0060] In some embodiments weights are assigned 1112 to each of the pairs of corresponding points. For example weights may be assigned using information from other sources such as an RGB camera or other sensors. . In an example, a weight related to measurement characteristics of the depth camera is computed and stored with each of the pairs of corresponding points. These weights may be used during the process of applying the error metric in order to improve the quality of the results.

[0061] In some embodiments pairs that include a point which is on or near a depth map boundary are rejected 1114. This helps to avoid errors where overlap between the two depth maps is only partial. Other criteria may also be used to reject pairs. For example, in some embodiments plane extraction is carried out as mentioned above with reference to FIG. 9 component 910. In that case, pairs which are on a plane may be rejected

[0062] In some embodiments, the source depth map is estimated or predicted from a dense 3D model of the scene being captured by the depth camera. In this case the method of FIG. 12 is followed. The dense 3D model of the scene comprises a 3D surface representation of the scene stored in GPU memory. Other ways of storing the 3D model may be used. For example, the dense 3D model may be stored as a linear array in slice-row-column order (more detail about this is given below), optionally with some padding so that slices and rows align certain memory

block sizes. Other ways of storing the 3D model may be used such as oct-trees, coarse-fine representations, mesh-based representations such as polygon meshes.

[0063] More detail about the case in which the dense 3D model is stored in

"slice-row-column" order on a GPU is now given. In this case, the model may be

5 stored as a linear array of memory locations used to represent a 3D volume. This is achieved by mapping each voxel to a memory array index using a linear pitched memory which provides fast, parallel access to the data stored on the GPU memory.

[0064] Surface normals of the sample points of the current depth map are

10 computed as described above by assessing 1200 neighboring points to the sample point. For the predicted source depth map, a surface normal prediction and a surface location prediction is computed 1202 from the dense 3D model for each predicted sample point. A predicted sample point is a point from the dense 3D

15 model which is at the same pixel location as a sample point from the current depth map. This is done by projecting a ray into the volume of the dense surface model. The ray is projected from an estimated camera position and orientation associated with the current depth map and into the 3D model through a point on a face of that 3D model which corresponds to a sample point in the current depth map. This applies in situations where the 3D model is stored as a volumetric representation.

20 In situations where the 3D model is stored using a mesh based representation then this representation is first projected to form a virtual depth image representation.

The ray may then be projected into that virtual depth image representation. A first visible surface along that ray is found by stepping along the ray and assessing a surface density function to find a first positive to negative zero crossing. The
25 associated sub pixel world point is found from an estimate of the intersection of the surface density function along the ray. In one example, the surface intersection point along a ray can be computed using a simple linear interpolation given trilinearly sampled points either side of the detected zero crossing to find the sub pixel world point at which a zero occurs. This sub pixel world point is taken as the
30 predicted surface location. To find the predicted surface normal at this location, finite differences of the surface density function gradient are found using tri-linear

interpolation. The process of calculating the surface normal prediction and surface location prediction may be implemented at the GPU with each ray being treated in parallel.

[0065] For each predicted sample point (obtained from the dense 3D model) a process 1204 is followed to identify a corresponding point in the current depth map. This is similar to the process 1104 of FIG. 11. The predicted sample point is projected 1206 onto a destination point in the destination depth map (current depth map). A search 1208 is then made for candidate corresponding points around the destination point which have compatible surface normals with the destination point. From those candidate corresponding points a point is selected 1110 according to a distance metric. For example, a pair of points are compatible if the points are within a specified Euclidean distance $e1$ of each other and the dot product between the surface normals of the pair is greater than a specified threshold $e2$. The parameters $e1$ and $e2$ may be user configurable or may be set during a manufacturing stage whereby the device is calibrated empirically for use in a particular setting.

[0066] In some cases weights are assigned 1212 to the pairs of corresponding points. In some embodiments pairs are rejected 1214 if they include at least one point which is on or near a depth map boundary. In an example a weight related to the measurement characteristics of the depth camera is stored with each of the pairs of corresponding points. These weights may be used during the process of applying the error metric in order to improve the quality of the results

[0067] Once pairs of corresponding points have been identified using the process of FIG. 11 or FIG. 12 for example, then an error metric is computed and minimized and the iterative process of FIG. 10 repeats.

[0068] In an example a point-to-plane error metric is computed 1300 for the pairs of corresponding points and this metric is optimized to obtain updated registration parameters. An example of this process is now described with reference to FIG. 13. This process is designed to be implemented using at least one parallel computing unit such as a GPU in order to obtain real time processing as now described.

[0069] Computing the point-to-plane error metric can be thought of as computing 1302 a sum of squared distances from each source point to a plane which contains the destination point and which is oriented perpendicular to the approximate surface normal of the destination point. The process seeks to optimize
 5 this metric to find an updated set of registration parameters. Solving this type of optimization problem is not straightforward and typically requires significant computational resources so making this type of process difficult to implement for real time applications. An example implementation using a parallel processing unit such as a GPU is now described which enables real-time processing.

10 [0070] The pairs of corresponding points may be scaled and translated 1304. This may improve the stability of the optimization process but is not essential.

[0071] For each pair of corresponding points a linear system comprising a plurality of simultaneous equations is formed 1306 on a parallel computing unit in order to optimize the error metric using numerical least squares optimization. Each
 15 linear system may be in the form of a 6 by 6 matrix. The plurality of matrices may be reduced to a single 6 by 6 matrix on the parallel computing unit. Because the frame rate is high (for example, 20 frames per second or more) then it is possible to make a small angle approximation for the angle (change in camera orientation) between any two successive frames. That is, because the frame rate is so high, the
 20 camera will only have moved a small amount between frames. By making this approximation the real time operation of the system is facilitated.

[0072] The single 6 by 6 matrix is passed to a CPU 1308 and solved to find updated registration parameters. The solution is scaled and translated 1310 back to reverse the scaling and translation step of 1304. The stability of the solution is
 25 checked 1312 and the updated registration parameters 1314 are output by the process.

[0073] In an example, the following point-to-plane error metric is used although this is not essential; other error metrics may also be used:

$$\arg \min_{\mathbf{T} \in \text{SE}_3} \sum_{\substack{\mathbf{u} \in \mathbf{U} \\ \rho_k(\mathbf{u}) \neq \text{null}}} \left((\mathbf{T} \mathbf{v}_k(\mathbf{u}) - \hat{\mathbf{v}}_{k-1, \rho_k(\mathbf{u})}^g) \cdot \hat{\mathbf{n}}_{k-1, \rho_k(\mathbf{u})}^g \right)^2 .$$

[0074] This error metric may be used to obtain a new transformation T_k .

More detail about the symbols used is now given. The current frame k of the depth camera D_k provides calibrated depth measurements $d = D_k(\mathbf{u})$ at image pixel

$\mathbf{u} = (x, y)$ in the image domain $\mathbf{u} \in U$. These measurements may be re-projected

5 into the camera's world space as $\mathbf{v}_k(\mathbf{u}) = (xd, yd, d, 1)$ (using homogeneous

coordinates). Since each frame from the depth sensor is a surface measurement on a regular grid, the system may also compute the corresponding normal vectors $\mathbf{n}_k(\mathbf{u})$

which are estimated by finite differences between neighbouring re-projected grid points. The SE3 transformation matrix maps the camera coordinate frame at time

10 k into the global frame g as $\mathbf{v}_k^g(\mathbf{u}) = \mathbf{T}_k \mathbf{v}_k(\mathbf{u})$. (The equivalent mapping of normal

vectors is $\mathbf{n}_k^g(\mathbf{u}) = \mathbf{R}_k \mathbf{n}_k(\mathbf{u})$). An estimate of the 3D model in the global coordinate

system at time k is denoted M_k which may be stored in a volumetric

representation, described herein. An incoming depth frame D_k is registered against the previous frame's estimate of the full 3D reconstructed model M_{k-1} , by ray-

15 casting into the previous frame's camera pose $\mathbf{T}_{k-1}$. This results in predicted image

$\hat{D}_{k-1}$ or equivalently a set of global model points $\hat{\mathbf{v}}_{k-1,i}^g$ and model normals $\hat{\mathbf{n}}_{k-1,i}^g$

where $i \in S$ is the corresponding index set. The symbol ρ_k in the above equation

for the point to plane error metric represents the protective data association

mapping between camera and model points at time k .

20 [0075] In an example, each pair of corresponding points that is identified by

the frame alignment engine 908 may be processed in parallel at the parallel

computing unit such as a GPU. Thus for each pair of corresponding points, a linear

system is computed which gives an arithmetic expression of a point-to-plane

constraint system. By making the small angle assumption the transformation $\mathbf{T}$

25 may be parametrised using a 3 vector of incremental rotations by a skew symmetric

matrix $\mathbf{R} \approx [\alpha, \beta, \gamma]_\times$ together with a 3 element translation vector $\mathbf{t}$. A linear system

is obtained by setting the first derivative of the linearised error metric to zero. This

point-to-plane constraint system expresses the optimization of the point-to-plane

error metric mentioned above. This computation occurs in parallel for each pair of

corresponding points at the parallel computing unit. In this way the error metric is applied to each of the identified corresponding points in parallel. The arithmetic expression for each pair of points is evaluated using a tree reduction process or other suitable method of evaluating an arithmetic expression. A tree reduction process is an evaluation strategy whereby an arithmetic expression is represented as a tree structure with nodes of the tree representing arithmetic operations and leaf nodes of the tree representing values. The expression is evaluated in an order according to the tree structure by passing results of evaluations along the branches of the tree. Together the results of the error metric optimization process from the parallel processes enable a single 6 by 6 matrix to be output.

[0076] FIG. 14 illustrates various components of an exemplary computing-based device 1404 which may be implemented as any form of a computing and/or electronic device, and in which embodiments of a real-time camera tracker may be implemented.

[0077] The computing-based device 1404 comprises one or more input interfaces 1402 arranged to receive and process input from one or more devices, such as user input devices (e.g. capture device 1406, a game controller 1405, a keyboard 1407, a mouse 1409). This user input may be used to control software applications or real-time camera tracking. For example, capture device 1406 may be a mobile depth camera arranged to capture depth maps of a scene. The computing-based device 1404 may be arranged to provide real-time tracking of that capture device 1406.

[0078] The computing-based device 1404 also comprises an output interface 1410 arranged to output display information to a display device 1408 which can be separate from or integral to the computing device 1404. The display information may provide a graphical user interface. In an example, the display device 1408 may also act as the user input device if it is a touch sensitive display device. The output interface 1410 may also output data to devices other than the display device, e.g. a locally connected printing device.

[0079] Computer executable instructions may be provided using any computer-readable media that is accessible by computing based device 1404.

Computer-readable media may include, for example, computer storage media such as memory 1412 and communications media. Computer storage media, such as

5 memory 1412, includes volatile and non-volatile, removable and non-removable media implemented in any method or technology for storage of information such as computer readable instructions, data structures, program modules or other data.

Computer storage media includes, but is not limited to, RAM, ROM, EPROM, EEPROM, flash memory or other memory technology, CD-ROM, digital versatile
10 disks (DVD) or other optical storage, magnetic cassettes, magnetic tape, magnetic disk storage or other magnetic storage devices, or any other non-transmission medium that can be used to store information for access by a computing device.

[0080] In contrast, communication media may embody computer readable instructions, data structures, program modules, or other data in a modulated data
15 signal, such as a carrier wave, or other transport mechanism. As defined herein, computer storage media does not include communication media. Although the computer storage media (memory 1412) is shown within the computing-based device 1404 it will be appreciated that the storage may be distributed or located remotely and accessed via a network or other communication link (e.g. using
20 communication interface 1413).

[0081] Computing-based device 1404 also comprises one or more processors 1400 which may be microprocessors, graphics processing units (GPUs), controllers or any other suitable type of processors for processing computing executable instructions to control the operation of the device in order to provide real-time
25 camera tracking. In some examples, for example where a system on a chip architecture is used, the processors 1400 may include one or more fixed function blocks (also referred to as accelerators) which implement a part of the method of real-time camera tracking in hardware (rather than software or firmware).

[0082] Platform software comprising an operating system 1414 or any other suitable platform software may be provided at the computing-based device to enable application software 1416 to be executed on the device. Other software than may be executed on the computing device 1404 comprises: frame alignment engine 1418 (see for example, FIG. 9 to 13 and description above), loop closure engine 1420, relocalization engine 1422. A data store 1424 is provided to store data such as previously received depth maps, registration parameters, user configurable parameters, other parameters, 3D models of scenes, game state information, game meta data, map data and other data.

[0083] The term 'computer' is used herein to refer to any device with processing capability such that it can execute instructions. Those skilled in the art will realize that such processing capabilities are incorporated into many different devices and therefore the term 'computer' includes PCs, servers, mobile telephones, personal digital assistants and many other devices.

[0084] The methods described herein may be performed by software in machine readable form on a tangible storage medium e.g. in the form of a computer program comprising computer program code means adapted to perform all the steps of any of the methods described herein when the program is run on a computer and where the computer program may be embodied on a computer readable medium.

Examples of tangible (or non-transitory) storage media include disks, thumb drives, memory etc and do not include propagated signals. The software can be suitable for execution on a parallel processor or a serial processor such that the method steps may be carried out in any suitable order, or simultaneously.

[0085] This acknowledges that software can be a valuable, separately tradable commodity. It is intended to encompass software, which runs on or controls “dumb” or standard hardware, to carry out the desired functions. It is also intended to encompass software which “describes” or defines the configuration of hardware, such as HDL (hardware description language) software, as is used for designing silicon chips, or for configuring universal programmable chips, to carry out desired functions.

[0086] Those skilled in the art will realize that storage devices utilized to store program instructions can be distributed across a network. For example, a remote computer may store an example of the process described as software. A local or terminal computer may access the remote computer and download a part or
5 all of the software to run the program. Alternatively, the local computer may download pieces of the software as needed, or execute some software instructions at the local terminal and some at the remote computer (or computer network). Those skilled in the art will also realize that by utilizing conventional techniques known to those skilled in the art that all, or a portion of the software instructions
10 may be carried out by a dedicated circuit, such as a DSP, programmable logic array, or the like.

[0087] Any range or device value given herein may be extended or altered without losing the effect sought, as will be apparent to the skilled person.

[0088] Although the subject matter has been described in language specific
15 to structural features and/or methodological acts, it is to be understood that the subject matter defined in the appended claims is not necessarily limited to the specific features or acts described above. Rather, the specific features and acts described above are disclosed as example forms of implementing the claims.

[0089] It will be understood that the benefits and advantages described above
20 may relate to one embodiment or may relate to several embodiments. The embodiments are not limited to those that solve any or all of the stated problems or those that have any or all of the stated benefits and advantages. It will further be understood that reference to 'an' item refers to one or more of those items.

[0090] The steps of the methods described herein may be carried out in any
25 suitable order, or simultaneously where appropriate. Additionally, individual blocks may be deleted from any of the methods without departing from the spirit and scope of the subject matter described herein. Aspects of any of the examples described above may be combined with aspects of any of the other examples described to form further examples without losing the effect sought.

[0091] The term 'comprising' is used herein to mean including the method blocks or elements identified, but that such blocks or elements do not comprise an exclusive list and a method or apparatus may contain additional blocks or elements.

[0092] It will be understood that the above description of a preferred
5 embodiment is given by way of example only and that various modifications may be made by those skilled in the art. The above specification, examples and data provide a complete description of the structure and use of exemplary embodiments of the invention. Although various embodiments of the invention have been
10 described above with a certain degree of particularity, or with reference to one or more individual embodiments, those skilled in the art could make numerous alterations to the disclosed embodiments without departing from the spirit or scope of this invention.

CLAIMS

1. A method of real time camera relocalization comprising:
 - receiving a sequence of depth map frames from a moving mobile depth camera each depth map frame comprising a plurality of image elements each having a depth value being related to a distance from the mobile depth camera to a surface in the scene captured by the mobile depth camera;
 - tracking the position and orientation of the mobile depth camera using the depth map frames and simultaneously forming a 3D model of the environment in which the mobile depth camera is moving using the depth map frames;
 - detecting a failure in the tracking of the position and orientation of the mobile depth camera; and
 - relocalizing the mobile depth camera by recomputing its position and orientation using a current depth map captured by the mobile depth camera.
2. A method as claimed in claim 1 wherein detecting a failure in the tracking comprises any of: comparing a change in a current and previous tracked position of the mobile depth camera with a threshold; detecting a failure in convergence of an iterative process used to track the position and orientation of the mobile depth camera; and comparing a current tracked position and orientation of the mobile depth camera with a position and orientation predicted using a model of motion of the mobile depth camera.
3. A method as claimed in claim 1 or claim 2 wherein relocalizing the mobile depth camera comprises finding a keyframe which is similar to the current depth map from a plurality of keyframes being depth map frames previously gathered by the mobile depth camera each having an associated camera position and orientation; and wherein finding a keyframe comprises any of: using a geometry estimation process to selecting a keyframe which has similar arrangement and locations of vertical and horizontal surfaces as the current frame; and searching the plurality of keyframes in an order related to a last known location of the mobile depth camera.

4. A method as claimed in claim 1 wherein relocalizing the mobile depth camera comprises: computing a track distribution of poses of the mobile depth camera being a predicted distribution of poses given known previous motion of the mobile depth camera; sampling a pose from the track distribution and using that sampled pose to align the current depth map and a previous depth map.

5. A method as claimed in claim 1 wherein relocalizing the mobile depth camera comprises: applying a random decision forest to patches from the current depth map and from a plurality of previous depth maps obtained from the 3D model of the environment to obtain histograms of textons being texture features of a depth map and selecting a previous depth map which is similar to the current depth map in terms of the histograms.

6. A real time camera relocalization system comprising:
an input arranged to receive a sequence of depth map frames from a moving mobile depth camera each depth map frame comprising a plurality of image elements each having a depth value being related to a distance from the mobile depth camera to a surface in the scene captured by the mobile depth camera;
a frame alignment engine arranged to track the position and orientation of the mobile depth camera using the depth map frames;
a 3D model formation system arranged to form a 3D model of the environment in which the mobile depth camera is moving using the depth map frames;
a relocalization engine arranged to detect a failure in the tracking of the position and orientation of the mobile depth camera; and to relocalize the mobile depth camera by recomputing its position and orientation using a current depth map captured by the mobile depth camera and at least one previous depth map obtained from the 3D model.


7. A system as claimed in claim 6 wherein the relocalization engine is arranged to apply a random decision forest to patches from the current depth map and from a plurality of previous depth maps obtained from the 3D model of the environment to obtain histograms of textons being texture features of a depth map and to select a previous depth map which is similar to the current depth map in terms of the histograms.

8. A system as claimed in claim 6 wherein the relocalization engine is arranged to find a keyframe which is similar to the current depth map from a plurality of keyframes being depth map frames previously gathered by the mobile depth camera each having an associated camera position and orientation.

9. A system as claimed in claim 8 wherein the relocalization engine is arranged to find a keyframe by any of: using a geometry estimator process to select a keyframe which has similar arrangement and locations of vertical and horizontal surfaces as the current frame; and searching the plurality of keyframes in an order related to a last known location of the mobile depth camera.

10. A game system comprising a mobile infra-red time-of-flight depth camera that uses structured light and a system as recited in claim 6 for tracking the mobile depth camera, the mobile depth camera and the real-time tracker being arranged to operate at at least 30 frames per second, the game system being arranged to influence the course of a game in relation to the tracking of the mobile depth camera.

1/14

**Fig. 1**






Fig. 2


3/14

**Fig. 3**


4/14

**Fig. 4**


5/14

**Fig. 5**


6/14

**Fig. 6**


7/14

**Fig. 7**


8/14

**Fig. 8**


9/14

**Fig. 9**

10/14

**Fig. 10**

11/14

**Fig. 11**

12/14






Fig. 12

13/14

**Fig. 13**

14/14

**Fig. 14**