

(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property Organization

International Bureau



(10) International Publication Number

WO 2016/182846 A1

(43) International Publication Date

17 November 2016 (17.11.2016)

WIPO | PCT

(51) International Patent Classification:

H04W 64/00 (2009.01) G01S 5/02 (2006.01)
H04W 4/02 (2009.01) G06T 7/00 (2006.01)

(21) International Application Number:

PCT/US2016/030956

(22) International Filing Date:

5 May 2016 (05.05.2016)

(25) Filing Language:

English

(26) Publication Language:

English

(30) Priority Data:

14/708,970 11 May 2015 (11.05.2015) US

(71) Applicant: GOOGLE INC. [US/US]; 1600 Amphitheatre Parkway, Mountain View, California 94043 (US).

(72) Inventors: WILLIAMS, Brian Patrick; 1600 Amphitheatre Parkway, Mountain View, California 94043 (US). TU, Laurent; Level 5, 48 Pirrama Road, Pyrmont, Australian Capital Territory NSW 2009 (AU). NERURKAR, Esha; 1600 Amphitheatre Parkway, Mountain View, California 94043 (US). LYNEN, Simon; ETH Zuerich, LEE J 317, Leonhardstrasse 21, 8092 Zuerich (CH).

(74) Agent: DAVIDSON, Ryan S.; Davidson Sheehan LLP, 700 Lavaca Street, Suite 1400-2323, Austin, Texas 78701 (US).

(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, BN, 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, IR, IS, JP, KE, KG, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SA, 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, ST, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, 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, KM, ML, MR, NE, SN, TD, TG).

Published:

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

(54) Title: PRIVACY-SENSITIVE QUERY FOR LOCALIZATION AREA DESCRIPTION FILE




FIG. 1

(57) Abstract: A computing system includes a datastore, a network interface, and a query module. The datastore stores a plurality of localization area description files. The network interface is to receive a request for a localization area description file from a mobile device, the request comprising a set of spatial features and at least one non-image location indicator. The query module includes a query interface to identify one or more candidate localization area description files based on one of the set of spatial features of the request and the at least one location indicator of the request, and includes a selection module to select a localization area description file from the candidate localization area description files based on the other of the set of spatial features of the request and the at least one location indicator. The query module is to provide the selected localization area description file to the mobile device.

WO 2016/182846 A1

PRIVACY-SENSITIVE QUERY FOR LOCALIZATION AREA DESCRIPTION FILE

Field of the Disclosure

The present disclosure relates generally to visual mapping systems and more particularly to the localization of mobile devices using visual information.

5 Background

Visual mapping systems rely on spatial features (also referred to as “visual features”) detected in imagery captured by a mobile device, as well as inertial information, to determine the current position and orientation of the mobile device in a three-dimensional (3D) space. Typically, the position and orientation are determined in the

10 context of a defined coordinate frame so as to facilitate various functionality that requires synchronization to a known, fixed reference frame, such as virtual reality (VR) functionality, augmented reality (AR) functionality, or gaming or other device-enabled interactions between multiple mobile devices. Simultaneous localization and mapping (SLAM) techniques enable a mobile device to map a previously unmapped area while concurrently learning its position and orientation within the area. Thus, when the mobile device returns to the same area, it may readily determine its current position and orientation within that area through detection of previously-observed spatial features in a process known as “localization.” However, when the mobile device is entering an area for the first time, the mobile device lacks these previously-15 detected localization cues. In conventional visual mapping systems the mobile device must “learn” the area through implementation of the visual mapping process – a process which takes considerable time and resources. To avoid the delay involved in performing the visual mapping process for a previously-unmapped area, conventional visual mapping systems instead may revert to detection of orientation or 20 position of the mobile device based on non-visual orientation input, such as global positioning system (GPS) information or location mapping via inertial sensor feedback. However, these non-visual mapping solutions can be unreliable (e.g., poor GPS reception indoors or in areas surrounded by tall obstructions), imprecise, and prone to error due to sensor and measurement drift.

Brief Description of the Drawings

The present disclosure may be better understood by, and its numerous features and advantages made apparent to, those skilled in the art by referencing the accompanying drawings. The use of the same reference symbols in different

5 drawings indicates similar or identical items.

FIG. 1 is a diagram illustrating a cloud-based visual mapping system for crowd-sourced creation and updating of area description files provided to mobile devices for localization in accordance with at least one embodiment of the present disclosure.

FIG. 2 is a block diagram illustrating an area description file server of the visual

10 mapping system of FIG. 1 in accordance with at least one embodiment of the present disclosure.

FIG. 3 is a block diagram illustrating a mobile device of the visual mapping system of FIG. 1 in accordance with at least one embodiment of the present disclosure.

FIG. 4 is a flow diagram illustrating a method for selecting one of a previously-

15 mapped localization process or a previously-unmapped localization process in accordance with at least one embodiment of the present disclosure.

FIG. 5 is a flow diagram illustrating a method for crowd-sourced generation of a

localization area description file in accordance with at least one embodiment of the present disclosure.

20 FIG. 6 is a flow diagram illustrating a method for localization of a mobile device using a localization area description file obtained from a remote area description file server in accordance with at least one embodiment of the present disclosure.

FIG. 7 is a flow diagram illustrating a method for crowd-sourced updating of a

localization area description file in accordance with at least one embodiment of the

25 present disclosure.

FIG. 8 is a flow diagram illustrating a method for evaluating a candidate spatial

feature for inclusion in a localization area description file in accordance with at least one embodiment of the present disclosure.

FIG. 9 is a flow diagram illustrating a method for evaluating a spatial feature for removal from a localization area description file in accordance with at least one embodiment of the present disclosure.

FIG. 10 is a diagram illustrating a two-level query interface for an LADF database of an area description file server for performing localization area description file queries in accordance with at least one embodiment of the present disclosure.

FIG. 11 is a flow diagram illustrating a method for a two-level query for a localization area description file in accordance with at least one embodiment of the present disclosure.

FIG. 12 is a diagram illustrating an alternative two-level query interface for an LADF database of an area description file server for performing localization area description file queries in accordance with at least one embodiment of the present disclosure.

DETAILED Description

The following description is intended to convey a thorough understanding of the present disclosure by providing a number of specific embodiments and details involving visual mapping systems. It is understood, however, that the present disclosure is not limited to these specific embodiments and details, which are examples only, and the scope of the disclosure is accordingly intended to be limited only by the following claims and equivalents thereof. It is further understood that one possessing ordinary skill in the art, in light of known systems and methods, would appreciate the use of the disclosure for its intended purposes and benefits in any number of alternative embodiments, depending upon specific design and other needs.

FIGs. 1-12 illustrate example systems and techniques for cloud-based creation, revision, and provision of localization area data files to mobile devices. In at least one embodiment, a visual mapping system comprises an area description file (ADF) server communicatively coupled to a plurality of mobile devices via one or more wired or wireless networks. When a mobile device enters a previously-unmapped area, the mobile device initiates an area learning process whereby the mobile device captures imagery of the area, detects spatial features using the captured imagery, and

transmits representations of the spatial features, their relative geometry, associated statistical data, and concurrently captured sensor data in an ADF to the ADF server.

The ADF server uses this ADF to generate a localization ADF (LADF), which represents a sparse point cloud of spatial features of the area. Note that the term

5 "file", such as used in "area description file" or as otherwise used herein, refers to any data structure used to associate data and other information fields, or any combination of such data structures. The LADF then may be downloaded to another mobile device in the area covered by the LADF. A mobile device receiving the LADF may capture imagery of the area, detect spatial features therein, and compare the
10 detected spatial features and their relative geometry to the spatial features and relative geometries represented by the LADF so as to localize the mobile device within the area using the reference coordinate frame presented by the LADF. In this manner, the area discovery and mapping of a previously-unmapped area may be efficiently crowd-sourced through passive collection and uploading of data by the
15 mobile devices.

In some embodiments, the ADF server receives ADF files from multiple mobile devices for the same area, or for adjacent areas, and the ADF server operates to merge these multiple ADF files into a merged ADF file (or merged cluster of ADF files), which in turn may be used to generate one or more LADFs for the area, or sub-

20 areas thereof. Further, as changes in the area are likely to occur, the ADF server may utilize crowd-sourced updating of an LADF generated for an area based on feedback from the mobile devices using the LADF for localization in the area. For example, in response to feedback indicating that some spatial features in the LADF are not being reliably observed by the mobile devices, the ADF server may remove
25 these spatial features from the LADF. Conversely, in response to feedback indicating that spatial features not included in the LADF are being reliably observed by mobile devices in the area, the ADF server may add these spatial features to the LADF.

Because the point clouds and associated data represented by ADFs and LADFs provide a visual representation of an area to a certain degree, the collection of ADFs
30 from mobile devices and the distribution of LADFs to mobile devices may have privacy implications. Accordingly, in some embodiments, the visual mapping system implements certain privacy controls. One such privacy control may include

configuring the mobile device to implement one or more privacy filter processes on the data of an ADF before the ADF is uploaded to the ADF server, and thus rendering the information represented by the uploaded version of the ADF unusable for the purposes of recreating useful visual content of the area. Another such privacy

5 control may include a two-stage LADF query for accessing LADFs for distribution to mobile devices. When a mobile device enters an area for which it does not have an LADF available, the mobile device captures imagery, detects spatial features from the captured imagery, and submits an LADF request to the ADF server, whereby the LADF request includes a set of the detected spatial features and an indicator of the

10 location of the mobile device (e.g., GPS coordinates or one or more wireless base station identifiers). The ADF server then may conduct a first query stage to identify a set of candidate LADFs that match one of the set of detected spatial features or the location indicator, and then select an LADF from the set of candidate LADFs based on the other of the detected spatial features or the location indicator. In this way, the

15 set of spatial identifiers submitted by the mobile device may serve as proof that the mobile device is in, or has access to, the identified area, and thus there is a significantly reduced risk of a privacy breach in supplying the LADF to the requesting mobile device.

FIG. 1 illustrates a visual mapping system 100 in accordance with at least one embodiment of the present disclosure. In the depicted example, the visual mapping system 100 comprises an ADF server 102 communicatively coupled to one or more mobile devices 104. As represented by cloud 106, the ADF server 102 comprises a computing system that is remote to the mobile device 104, and is coupled to the mobile device 104 via one or more wired or wireless networks, such as via a wireless local area network (WLAN), a cellular data network, the Internet, or a combination thereof. Although described herein in the example context of a single server, in other embodiments the ADF server 102 may be implemented as a computing system comprising a cluster of servers. An example implementation of the ADF server 102 is described in greater detail below with reference to FIG. 2.

30 The mobile device 104 can include any of a variety of portable electronic devices operated by one or more users 110, such as head mounted display (HMD), a tablet computer, computing-enabled cellular phone (e.g., a “smartphone”), a notebook

computer, a personal digital assistant (PDA), a gaming console system, a drone, and the like. In the depicted example, the mobile device 104 includes a housing 112 having a surface 114 opposite another surface 116, whereby the housing 112 typically is oriented with respect to the user 110 such that the user faces the surface 114 of the housing 112. Further, in the depicted implementation, the mobile device 104 includes a display 118 disposed at the surface 116 for presenting visual information to the user 110. Accordingly, for ease of reference, the surface 116 is referred to herein as the “forward-facing” surface and the surface 114 is referred to herein as the “user-facing” surface as a reflection of this example orientation, 5 although the orientation of these surfaces is not limited by these relational designations. 10

The mobile device 104 further includes a plurality of sensors to obtain information regarding an area 122 in which the mobile device 104 is currently located. The mobile device 104 obtains visual information (imagery) for the area 122 via one or 15 more imaging sensors, such as imaging sensors 124, 126, disposed at, for example, the forward-facing surface 116. The imaging sensors 124, 126 can be positioned and oriented on the forward-facing surface 116 such that their respective fields of view overlap starting at a specified distance from the mobile device 104, thereby enabling, via multiview analysis, depth sensing of objects in the area 122 that are 20 positioned in the region of the overlapping fields of view. Alternatively, a depth sensor 130 disposed on the surface 116 may be used to provide depth information for the objects in the area.

The depth sensor 130, in one embodiment, uses a modulated light projector to project modulated light patterns from the forward-facing surface 116 into the area 25 122, and uses one or both of imaging sensors 124, 126 to capture reflections of the modulated light patterns as they reflect back from objects in the area 122. These modulated light patterns can be either spatially-modulated light patterns or temporally-modulated light patterns. The captured reflections of a modulated light flash are referred to herein as “depth images” or “depth imagery.” The depth sensor 30 120 then may calculate the depths of the objects, that is, the distances of the objects from the mobile device 104, based on the analysis of the depth imagery. The resulting depth data obtained from the depth sensor 130 may be used to calibrate or

otherwise augment depth information obtained from multiview analysis (e.g., stereoscopic analysis) of the image data captured by the imaging sensors 124, 126. Alternatively, the depth data from the depth sensor 130 may be used in place of depth information obtained from multiview analysis. A more detailed example of an 5 electronic system of the mobile device 104 is described in greater detail below with reference to FIG. 3.

In operation, the mobile device 104 captures imagery of the area 122 via one or both of the imaging sensors 124, 126, modifies or otherwise processes the captured imagery, and provides the processed captured imagery for display on the display 10 108. The processing of the captured imagery can include, for example, spatial or chromatic filtering, addition of an augmented reality (AR) overlay, conversion of real-life content of the imagery to corresponding virtual reality (VR) content, and the like. For the provision of this AR or VR functionality, the mobile device 104 relies on an accurate determination of its current six degree of freedom 15 (6DOF) orientation and position (herein collectively referred to as "pose") relative to a specified coordinate frame. Thus, in some embodiments, the mobile device 104 also uses imagery captured by the imaging sensors 124, 126 and non-image sensor data (e.g., inertial sensor data) to determine a relative position/orientation of the mobile device 104, that is, a position/orientation relative to the area 122, using one or more 20 of a simultaneous location and mapping (SLAM) process, a visual odometry process, or other visual mapping process.

Conventionally, when a mobile device enters a new area (that is, an area that the particular mobile device has not previously entered), the mobile device either must complete an area learning process (e.g., a SLAM process) to learn the new area, or 25 rely on non-visual cues (e.g., GPS coordinates) to provide some indication of the current pose of the mobile device relative to a coordinate frame. In contrast, the visual mapping system 100 leverages crowdsourcing provided through the prior explorations by other mobile devices in the area to allow a mobile device new to the area to quickly localize to a coordinate frame without extensive visual mapping of the 30 area. To this end, mobile devices 104 connected to the ADF server 102 operate to collect and upload ADF 132 to the ADF server 102 as the mobile devices 104 move about the area 122. As described in greater detail below, the ADF 132 includes a set

of spatial features detected by the mobile device 104, the relative geometry of the spatial features, various statistical metadata regarding the spatial features that describe the spatial features in ways substantially independent of the mobile device's pose, the lighting conditions, or other temporary environmental conditions or device-specific conditions, as well as sensor data collected from an inertial management unit (IMU) or other non-image sensors at the time of capture of the imagery associated with the spatial features. The ADF server 102 filters and merges the ADFs from multiple mobile devices 104 for the same area (or adjacent areas) based on the information contained in the ADF 132, and from the resulting merged ADFs 5 generates one or more LADFs 134, each of which represents a sparse point cloud of spatial features for the area or a corresponding sub-area thereof.

10

Thus, when another mobile device 104 subsequently enters a new area, the mobile device 104 may query the ADF server 102 for an LADF that is associated with the area. Responsive to this query, the ADF server 102 may provide an LADF 134 to the 15 requesting mobile device 104. The mobile device 104 may capture imagery of the area 122, detect certain spatial features contained therein, and compare the detected spatial features and their relative geometry with the spatial features and associated relative geometry represented by the LADF 134, and localize its current pose (that is, either determine its current pose or correct a previously determined pose that has 20 drifted) based on the comparison of spatial features, such as through the application of a loop closure algorithm. The mobile device 104 further may provide feedback to the ADF server 102 based on the mobile device's usage of the LADF 134, such as feedback data indicating that a spatial feature of the LADF 134 was not observed by the mobile device 104, or feedback data indicating that a spatial feature not in the 25 LADF 134 was observed by the mobile device 104.

The crowd-sourced and cloud-based creation, refinement, and distribution of LADFs 134 permits mobile devices 104 to quickly and efficiently localize when entering a new area, and thus avoid the need to perform a time-consuming area learning process or the need to use the less precise and often unreliable measurements from 30 a GPS sensor or inertial sensor. Moreover, as the spatial features of the LADF 134 may be referenced to a specific coordinate frame, multiple mobile devices 104 using the LADF 134 may localize their current poses to a common coordinate frame,

thereby facilitating more efficient and accurate interaction between the mobile devices 104 for functionality that relies on device pose information, such as multiplayer games, shared AR or VR functionality, and the like.

FIG. 2 illustrates an example implementation of the ADF server 102 of FIG. 1 in accordance with at least one embodiment of the present disclosure. In the depicted example, the ADF server 102 comprises a computing system having a network interface 202, an ADF datastore 204, an LADF datastore 206, a feature score datastore 208, a georeference datastore 210, a merge module 212, a georeference module 214, a spatial feature filter module 216, a localization generation module 218, and a query module 220. Although illustrated as separate datastores, one or more of the datastores 204, 206, 208, 210 may be implemented together as a single datastore.

The modules 212, 214, 216, 218, 220 may be implemented as hard-coded logic (e.g., an application specific integrated circuit or programmable logic), one or more processors 222 executing software instructions 224 stored in a memory 226 or other storage device, or a combination thereof. Further, although depicted as a single server for ease of illustration, the ADF server 102 instead may be implemented as a computing system comprising a plurality of servers. For example, the functionality of the modules 212, 214, 216, and 218 may be implemented at one server, and the functionality of the query module 220 and LADF datastore 206 may be implemented at another server.

As a general overview, the merger module 212 operates to receive, via the network interface 202, ADFs 132 from one or more mobile devices 232 (one embodiment of the mobile device 104 of FIG. 1) that have entered an area for which an LADF has 25 not yet been compiled. The merger module 212 merges the one or more ADFs 132 from the mobile devices 232 to generate merged ADF data for storage in the ADF datastore 204. The georeference module 214 may utilize georeference information from the georeference datastore 210 (which may contain, for example, imagery and associated georeference location from the Street View tool provided by Google Inc.) 30 to provide a geographic reference to the merged ADF. The localization generation module 218 generates one or more LADFs from the resulting merged ADF data and to store and index the one or more LADFs in the LADF datastore 206 for subsequent

retrieval. The query module 220 receives an LADF request 236 from a mobile device 234 (one embodiment of the mobile device 104) via the network interface 202, searches the LADF datastore 206 for the LADF 134 corresponding to the LADF request 236, and provides the LADF 134 to the mobile device 234. The mobile 5 device 234 uses the LADF 134 for localization of the mobile device 234. During this localization process, the mobile device 234 may provide LADF feedback 238 on the LADF 134 to the ADF server 102 via the network interface 202. The ADF server 102 may use the LADF feedback 238 to adjust feature scores of spatial features represented therein in the feature score datastore 208, and based on the feature 10 scores refine the LADF 134 by adding or removing spatial features. These operations are described in greater detail below with reference to FIGs. 4-12.

As noted, mobile devices 232 and 234 represent embodiments of the mobile device 104. Due to their respective operational requirements, a mobile device 232 performing a localization process using an LADF may not require the same 15 capabilities as a mobile device 234 performing the ADF upload process. To illustrate, the mobile device 234 may utilize the depth sensor 120 or a stereo camera configuration to facilitate SLAM operations in association with the process of generating and uploading an ADF, whereas the mobile device 232 typically will not require a depth sensor 120 and only requires a monocular camera configuration to 20 facilitate the localization process.

FIG. 3 illustrates an example processing system 300 implemented by the mobile device 104 in accordance with at least one embodiment of the present disclosure. The processing system 300 includes the display 118, the imaging sensors 124, 126, and the depth sensor 130. The processing system 300 further includes a graphics 25 processing unit (GPU 302), frame buffers 303 and 305, an application processor 304, a display controller 306, a system memory 308, a set 310 of non-image sensors, and a user interface 312. The user interface 312 includes one or more components manipulated by a user to provide user input to the mobile device 104, such as a touchscreen 314, a mouse, a keyboard, a microphone 316, various buttons or 30 switches, and various haptic actuators 318. The set 310 of non-image sensors can include any of a variety of sensors used to provide non-image context or state of the mobile device 104. Examples of such sensors include an inertial management unit

(IMU) 320 comprising one or more of a gyroscope 321, a magnetometer 322, and an accelerometer 323. The non-image sensors further can include, for example, an ambient light sensor 326 as well as various wireless reception or transmission based sensors, such as a GPS sensor 328, a wireless local area network (WLAN) interface 330, a cellular interface 332, a peer-to-peer (P2P) wireless interface 334, and a near field communications (NFC) interface 336. The non-image sensors also can include user input components of the user interface 312, such as the touchscreen 314 or the microphone 316.

The mobile device 104 further has access to various datastores 338 storing information or metadata used in conjunction with its image processing, location mapping, and location-utilization processes. The datastores 338 can include a spatial feature datastore to store metadata for 2D or 3D spatial features identified from imagery captured by the imaging sensors of the mobile device 104, a SLAM datastore that stores SLAM-based information, such as mapping information for sub-areas of the area 122 (FIG. 1) already explored by the mobile device 104, and an AR datastore that stores AR overlay information or VR information, such as CAD-based representations of the relative locations of objects of interest in the area 122. The datastores 338 may be implemented at one or more storage components of the mobile device 104, such as on a hard drive, solid state memory, or removable storage medium (not shown).

In operation, the imaging sensors 124, 126 capture imagery of an area and buffer the captured imagery in the frame buffers 303, 305. For imagery to be displayed in original or modified form, the GPU 302 processes the captured imagery for display (e.g., by rendering an AR overlay), and the display controller 306 controls the display 118 to display the processed imagery. Further, as described herein, the mobile device 104 operates to upload ADFs 134 for previously-unmapped locations, as well as to download LADFs for previously-mapped locations and use a downloaded LADF to facilitate localization of the mobile device 104. To this end, one or more software programs may be stored in the system memory 308 or other non-transitory computer 30 readable medium and executed by one or both of the application processor 304 and the GPU 302 to provide this ADF generation and LADF utilization functionality. For ease of illustration, the one or more software programs are depicted in FIG. 3 as an

ADF generation program 342 and an LADF processing program 344. These programs may be implemented as a thread or other process in an operating system (OS) executed by the processing system 300, as threads, processes or subroutines of the same software application, or as separately-executed software applications.

5 Further, in some embodiments, some or all of the functionality of the programs 342, 344 described herein may be implemented via an ASIC, programmable logic, or other hard-coded logic.

The ADF generation program 342 can include, for example, a privacy filter module 346, a spatial feature detection module 348, and an ADF assembly module 350. As

10 depicted in FIG. 3, the privacy filter module 346 may comprise one or more image content filters, such as a text filter module 352 and a facial filter module, as well as a downampler module 356. The LADF processing program 344 can include, for example, a spatial feature detection module 358 (which may be the spatial feature detection module 348), a request module 360, a localization module 362, and a 15 feedback module 364. The operations of the processing system 300, including the operations of the ADF generation program 342 and the LADF processing program 344, are described in detail below.

As described above, when the mobile device 104 is in an unmapped area (that is, an area for which the ADF server 102 does not have an LADF available), the mobile

20 device 104 may operate in an ADF generation mode whereby the mobile device 104 generates an ADF representing spatial features detected in the unmapped area and upload the ADF to the ADF server 102. However, if the area has already been mapped and the ADF server 102 has an LADF for the area, the mobile device 104 instead may operate in an LADF localization mode whereby the mobile device 104 25 obtains the LADF for the area from the ADF server 102 and uses the sparse point cloud represented by the LADF to perform a localization process so as to localize the mobile device 104 to the area.

FIG. 4 illustrates an example method 400 for selecting between these operational modes in accordance with at least one embodiment. For ease of description, the 30 method 400 is described in the example context of the processing system 300 of FIG. 3. The method 400 initiates at block 402, whereby the mobile device 104 determines that it has moved into an area that has not been previously mapped by the mobile

device 104. In response, the mobile device 104 initiates a localization process or motion tracking process.

At block 404, the mobile device 104 queries the ADF server 102 to determine whether an LADF is available for the area. As described below with reference to

5 FIGs. 10-12, the ADF server 102 may implement a two-stage query process whereby the mobile device 104 provides both a set of spatial features from imagery captured by the mobile device in the area and one or more location indicators (e.g., GPS coordinates or base station identifiers of WLAN or cell tower base stations detected by the mobile device 104) and the ADF server 102 queries the LADF datastore 206 10 to identify the corresponding LADF.

In the event that an LADF is not available from the ADF server 102, the area is considered to be an unmapped location and thus at block 406 the mobile device 104 and the ADF server 102 coordinate to perform an ADF/LADF generation process for the unmapped location. This process is described in greater detail below with

15 reference to FIG. 5. In the event that an LADF is available from the ADF server 102, the area is considered to be a mapped location and thus at block 408 the mobile device 104 and the ADF server 102 coordinate to perform an LADF localization and update process, which is described in greater detail below with reference to FIG. 6.

FIG. 5 illustrates an example method 500 for implementing the ADF/LADF generation process in accordance with at least one embodiment. As noted above, the ADF/LADF generation process is performed responsive to the determination that an LADF is not available for an area newly encountered by the mobile device 104.

Accordingly, after receiving user consent for the passive collection of area description data, the method 500 initiates at block 502 with the initiation of an area learning 25 process using the sensors of the mobile device 104. To perform the area learning process, the mobile device 104 captures imagery 372 (FIG. 3) via imaging sensors 124, 126 and captures corresponding depth information 373 via the depth sensor 130 as the mobile device 104 moves through the area. The mobile device 104 may employ, for example, an interactive game to induce the user of the mobile device 104 30 to explore the area as the imagery 372 and depth information 373 is captured by the mobile device 104.

Concurrently, the ADF assembly module 350 captures non-image sensor data 374 (FIG. 3) from one or more of the non-image sensors of the set 310. To illustrate, the ADF assembly module 350 may capture sensor data from the accelerometer 323 during the capture of the imagery 372, and thus represent the orientation of the

5 mobile device 104 with respect to gravity at the time of the imagery capture.

Similarly, sensor data captured from the gyroscope 321 may be used to determine the direction of travel of the mobile device 104 with respect to the visual features represented in the captured imagery. Further, the non-image sensor data 374 may include sensor data that may operate as location indicators of the mobile device 104.

10 These location indicators may be georeferenced location indicators, such as latitude/longitude coordinates represented in the sensor data provided by the GPS sensor 328. Alternatively, these location indicators may be inferred location indicators. To illustrate, as WLAN base stations and cell tower base stations are presumed to be fixed, the detection of a WLAN base station or a cell tower base 15 station serves as an indication that the mobile device is in proximity to the detected base station, and thus the base station identifier (BSID), media access control (MAC) address, or other identifier of the base station can serve as an inferred location indicator for the mobile device.

As described below, the captured imagery 372 and depth information 373 is used to 20 determine a point cloud of spatial features which ultimately is uploaded to the ADF server 102. While this point cloud is not image data itself, depending on the density of the point cloud and other conditions, it is possible that an unmodified point cloud could be used to reconstruct certain visual content originally present in the captured imagery 372. For example, if the mobile device 104 were placed particularly close to 25 a document, it is possible that the point cloud determined from the imagery captured of the document could be used to reproduce the text of the document. Accordingly, to prevent disclosure of unintended visual content, the privacy filter module 346 may implement at least two privacy controls.

The first privacy control is implemented at block 504, wherein the privacy filter 30 module 346 performs one or more content filter processes on the captured imagery 372 to remove image content from regions that may have privacy implications. For example, at block 504 the text filter module 352 may perform a text filter process

whereby each image captured at block 502 is scanned using one or more well-known text recognition algorithms to determine whether there are any regions in the image that potentially represent text content. For each region detected as a potential text region, the text filter module 352 may obfuscate or delete this region by, for example,

5 performing a blending operation with pixels in the region or in adjacent regions, by replacing the pixel values in the region with the same default pixel region or otherwise delete the pixel values in the region, and the like. Likewise, at block 504 the facial filter module 354 may implement a facial filter process whereby each captured image is scanned using one or more well-known facial recognition

10 algorithms to determine whether there are any regions in the image that potentially represent a human face, and filter the image to remove the image content from each region so identified. In this manner, the imagery is pre-filtered to remove potentially sensitive visual content prior to spatial feature detection, and thus the resulting point cloud of spatial features cannot be used to reconstruct the potentially sensitive visual

15 content.

At block 506, the spatial feature detection module 348 analyzes the filtered imagery to detect spatial features contained therein. Any of a variety of spatial feature extraction algorithms may be used, such as, for example, a Scale-Invariant Feature Transform (SIFT) algorithm, a Speeded-Up Robust Features (SURF) algorithm, a

20 Gray Level Patch algorithm, a Gradient Location and Orientation Histogram (GLOH) algorithm, a Zernike Moment algorithm, a Binary Robust Independent Elementary Features (BRIEF) algorithm, Oriented BRIEF (ORB) algorithm, Binary Robust Invariant Scalable Keypoints (BRISK) algorithm, Difference of Gaussians (DOG) algorithm, Fast Retina Keypoint (FREAK) algorithm, and the like. The spatial feature

25 detection module 348 provides these detected spatial features as spatial feature data 376. To illustrate, the FREAK algorithm provides for the comparison of pairs of pixels within an image patch, with the output of each comparison being a “0” or a “1” based on which pixel is lighter or darker of the pair. In the FREAK algorithm, 512 such comparisons are computed over the image patch, and the result is a spatial feature

30 descriptor that is a binary string of length 512 for a corresponding spatial feature representing the image patch, as well as an (x,y,z) vector identifying the position of the spatial feature in the 3D reference frame.

Further, the spatial feature detection module 348 determines statistical metadata 378 for each detected point. The statistical metadata 378 describes the corresponding spatial feature in a manner that is substantially independent of the particular viewing angle of the mobile device 104 or of the ambient lighting when the corresponding imagery was captured. For example, the statistical metadata 378 may comprise values representing the mean and standard deviation of the brightness of pixels representing the spatial feature, brightness gradients (or other visual characteristic gradients) in one or more directions, and the like.

Although the privacy filter module 346 scrubbed the imagery of detected text and

10 facial content, and thus the original point cloud of spatial features represented by the spatial feature data 376 is substantially devoid of textual and facial content, the original point cloud may still have sufficient spatial feature density to permit some level of reconstruction of visual appearance of the area from the original point cloud. Thus, as a second privacy control, at block 508 the downsample module 356 of the 15 privacy filter module 346 downsamples the original point cloud so as to generate a filtered point cloud 380 (FIG. 3) of spatial features that contains only a selected subset of the spatial features of the original point cloud. This has the benefit of both reducing the likelihood that the visual appearance of an area may be reconstructed from the resulting point cloud, as well as reducing the amount of data needed to 20 represent the resulting filtered point cloud 380. The downsampling process may be controlled by one or more downsampling criteria. For example, in some embodiments, the downsampling process may include the random selection of spatial features for exclusion from the filtered point cloud, the removal of every Xth spatial feature (X being an integer greater than 2), or the selection of no more than a 25 maximum number of spatial features for inclusion in the filtered point cloud. As another example, the downsampling process may be controlled by a maximum spatial feature density criterion that specifies a maximum number of spatial features per cubic unit (that is, unit volume). For example, the maximum spatial feature density criterion may specify that the filtered point cloud 380 contain at most one 30 spatial feature per cubic foot, and thus downsample the original point cloud so that no more than one spatial feature per cubic foot is represented in the resulting filtered point cloud.

At block 510, the ADF assembly module 350 uses the filtered point cloud 380, statistical metadata 378, and the non-image sensor data 374 to generate an ADF 132. In one embodiment, the ADF 132 comprises a file or other data structure that represents the filtered point cloud 380 as a list or other set of multidimensional coordinates, along with a field for each multidimensional coordinate to store the statistical metadata 378 for the represented spatial feature. To illustrate, each multidimensional coordinate may comprise an (x,y,z) floating point vector that represents a 3D position of a corresponding spatial feature in the (X,Y,Z) coordinate frame of the mobile device 104. The data structure further may include one or more fields to store the non-image sensor data 374, as well as a field to store a unique identifier (UID) assigned to the ADF 132 by the ADF server 102, as described below.

At block 512, the ADF assembly module 350 provides the ADF 132 to the appropriate network interface (e.g., the cellular interface 332 or the WLAN interface 330) for transmission to the ADF server 102. In some embodiments, the ADF assembly module 350 signals a request to upload an ADF to the ADF server 102, and the ADF server 102 responds with a UID assigned to the ADF, which the ADF assembly module 350 inserts into the appropriate field in the ADF 132 before it is uploaded. Further, while the process of blocks 502-512 has been described as a sequential process whereby the ADF 132 is completed before it is uploaded, in some embodiments, the ADF 132 is generated and uploaded in an iterative process whereby the ADF 132 is generated as a sequence of ADF blocks, each ADF block containing a subset of the point cloud generated from a portion of the captured imagery, and each ADF block being marked with the UID assigned to the ADF and uploaded to the ADF server 102 as it is generated. The ADF server 102 thus may store these ADF blocks individually, or may combine the ADF blocks denoted by the same UID into a single ADF 132.

The ADF 132 uploaded by the mobile device 104 represents certain visual features present in the area in which the mobile device 104 is located. However, through the privacy controls implemented by the mobile device 104 as described above with reference to blocks 504 and 508, the information contained in the ADF 132 that is provided to the ADF server 102 is substantially devoid of any content that may have privacy implications. Rather, while the resulting ADF 132 contains a sparse point

cloud that may describe various edges, corners, and other visual features of the area so to a degree sufficient to allow another mobile device to subsequently localize based on these described visual features, the sparse point cloud contains insufficient information to support replication of visual appearances of the area in a way

5 meaningful to human perception. As such, the ADF server 102 should not come into possession of any information from the mobile device 104 that has a measurable potential to reveal potentially sensitive information regarding a mapped area.

The role of the ADF server 102 in the ADF/LADF generation process initiates at block 514 with the receipt of the ADF 132 by the merge module 212 via the network

10 interface 202 from the mobile device 104. Upon receipt of the ADF 132, the merge module 212 temporarily indexes and stores the ADF 132 in the ADF datastore 204. The ADF 132 may be indexed for temporary storage based on the UID assigned to the ADF 132, based on the spatial features represented in the point cloud represented by the ADF 132, based on the location indicator(s) included with the 15 ADF 132, and the like. In at least one embodiment, each ADF 132 is stored and indexed in the ADF datastore 204 using the two-layer query approach described below with reference to FIGs. 10-12.

In some instances, multiple mobile devices may have uploaded a corresponding ADF for an area or for adjacent areas. The ADF server 102 thus operates at block 516 to

20 merge these co-located ADFs before processing the resulting merged ADFs into one or more LADFs. The merge process of block 516 may be initiated in response to any of a variety of triggers. For example, the merge process may be triggered by the lapse of a specified amount of time, by the receipt of a specified number of ADFs from mobile devices, by the request for an LADF from a mobile device, and the like.

25 In response to such a trigger, at block 518 the merge module 212 signals the query module 220 to query the ADF datastore 204 to identify whether there are any “nearby” ADFs – that is, ADFs that cover the same area or an adjacent area. In some embodiments, nearby ADFs may be identified based on comparisons of the location indicators associated with ADFs in the ADF datastore 240. For example, 30 assuming the upload of an ADF 132 from the mobile device 104 triggered the merge process and the uploaded ADF 132 includes one or more WLAN MAC addresses as location indicators, then those ADFs in the ADF datastore 240 identified by the query

module 220 as having the same WLAN MAC address as supplied with the uploaded ADF, or with WLAN MAC addresses of WLAN base stations known to be adjacent to the area covered by a WLAN base station identified by the uploaded ADF, are identified as nearby ADFs. As another example, the uploaded ADF may be provided 5 with GPS coordinates as a location indicator, and the query module 220 may identify ADFs in the ADF datastore 240 as being nearby based on their corresponding GPS coordinates. In other embodiments, nearby ADFs may be identified based on a comparison of the spatial features represented in the uploaded ADF with the spatial features represented in ADFs stored in the ADF datastore 240. Those ADFs having 10 a sufficiently overlapping set of spatial features thus may be identified as representing the same area or a nearby area, and thus identified as nearby ADFs. Still further, as described below with reference to FIGs. 10-12, the ADF server 102 may implement a two-layer query for one or both of the ADF datastore 204 and the LADF datastore 206 such that the each stored ADF/LADF is indexed for storage 15 based on both its spatial feature set and one or more location indicators, and nearby ADFs are identified by performing a two-layer indexing process using both the spatial feature set of the uploaded ADF as well as one or more location indicators of the uploaded ADF.

The set of one or more nearby, or co-located, ADFs identified by the query module 20 220 are referred to herein as an “ADF cluster.” With an ADF cluster identified, at block 520 the merge module 212 operates to perform one or more deduplication processes to remove duplicate spatial features represented in the ADFs of the ADF cluster. Further, in the event that the ADF cluster represents an excessively large area, or contains an excessively large number of spatial features, the merge module 25 212 may split the ADF cluster into one or more smaller ADF clusters, each of which may be processed as described herein. At block 522, the merge module 212 employs one or more well-known loop closure algorithms to determine the relative alignment between the ADFs of the ADF cluster. At block 524, the merge module 212 analyzes the ADF cluster to selectively remove spatial features which are not 30 reliably observed in most or all of the ADFs of the ADF cluster, as identified using one or more stability criteria. This serves both the purpose of eliminating spatial features which are likely not representative of permanent visual features of the area (and thus unreliable sources for localization), as well as reducing the total number of spatial

features in the combined point cloud represented by the ADF cluster. In selecting the spatial features to remove, the merge module 212 may employ a scoring system based on a stability criterion, such as by giving spatial features scores based on the frequency that they are observed within the ADFs of the ADF cluster. Further, as 5 newer ADFs are more likely to represent the current state of the area, the merge module 212 may give inclusion preference to spatial features in newer ADFs.

In some instances, the non-image sensor data 374 supplied with an ADF may permit the ADF to be georeferenced. For example, the non-image sensor data 374 may include GPS coordinates and an indicator of the georeferenced pose of the mobile

10 device 104 at the time of the detection of the spatial features in the ADF. With the relative alignment of the ADFs of the ADF cluster determined at block 522, if one ADF in the ADF cluster is georeferenced, then at block 526 the other ADFs in the ADF cluster may be georeferenced based on application of a coordinate frame translation using the relative alignment and geolocation information of the

15 georeferenced ADF. Further, in the event that georeferenced visual reference data is available, such as via the Street View tool provided by Google Inc., the ADFs of the ADF cluster may be georeferenced using this georeferenced visual reference data.

With the ADFs of the ADF cluster merged, at block 528 the localization generation module 218 uses the merged ADFs to generate one or more LADFs for the areas

20 represented by the merged ADFs. In the event that the data size of the merged ADFs are sufficiently small, or cover a sufficiently small area, the merged ADFs may be stored and indexed as a single LADF. However, in the event that the data size of the merged ADFs is over a threshold, the merged ADFs cover too large of an area, or the area contains multiple separate sub-areas (as identified via walls, partitions,

25 doors, and windows in the point clouds of the ADFs), the localization generation module 218 may spatially partition the merged ADFs to generate multiple LADFs, each LADF covering a different sub-area. In such instances, the localization generation module 218 may attempt to identify logical partition lines, such as by identifying multiple rooms within the area based on the detection of walls or other

30 room dividers within the point cloud represented by the merged ADFs, and create a separate LADF for each identified room, which has the additional benefit from a privacy perspective of limiting the scope of the LADF to a single room.

For each LADF generated, at block 530 the localization generation module 218 provides the LADF to the LADF datastore 206 for indexing and storage. As noted above and as described in detail below, in some embodiments the LADF datastore 206 employs a two-layer index for each LADF, whereby each LADF is indexed both 5 by one or more location indicia and by the spatial features represented by the LADF. Accordingly, when storing an LADF in the LADF datastore 206, the LADF datastore 206 stores the LADF in a datastore entry indexed by one of its location indicator(s) or its spatial feature set (and relative geometries therein), and the other of the location indicator(s) or the spatial feature set is used to select between multiple LADFs 10 similarly indexed.

With the conclusion of the process at block 530 of method 500, the ADF server 102 has generated one or more LADFs for a previously unmapped area using one or more ADFs uploaded from one or more mobile devices 104 that have had the opportunity to explore the area. Thus, the ADF server 102 is ready to supply a LADF 15 to any other mobile device that encounters this area for the first time, as described by FIG. 6 below.

FIG. 6 illustrates an example method 600 for performing the LADF localization and update process represented by block 408 of FIG. 4 in accordance with at least one embodiment. The method 600 initiates at block 602 when the mobile device 104 20 enters an area that the mobile device 104 has not previously encountered and mapped. In response to this determination, the mobile device initiates a spatial feature detection process for the area. To this end, the mobile device 104 triggers the imaging sensors 124, 126 to begin the capture of imagery 382 (FIG. 3) of the area and triggers the depth sensor 130 to begin the capture of depth information for 25 the area, and from the imagery 382 and depth information the spatial feature detection module 358 detects an initial set of spatial features representative of the area. At block 604, the request module 360 determines one or more location indicators for the location of the mobile device 104 at the time of spatial feature detection in the area. As noted above, these location indicators may be specific 30 geographic location indicators, such as GPS coordinates obtained from the GPS sensor 328, inferential location indicators, such as WLAN MAC addresses, WLAN BSIDs, or cell tower BSIDs, or a combination thereof. At block 606, the request

module 360 uses the initial spatial feature set and the one or more location indicators to generate an LADF request 236 and transmit the LADF request 236 to initiate an inquiry into the availability of an LADF for the area.

In response to receipt of the LADF request 236 at the query module 220 of the ADF 5 server 102 via the network interface 202, at block 608 the query module 220 queries the LADF datastore 206 to identify a suitable LADF for the area for use by the mobile device 104. In at least one embodiment, the query module 220 and the LADF datastore 206 employ a two-layer indexing scheme to identify a suitable LADF based on both the initial spatial feature set and the location indicators, as described in detail 10 below with reference to FIGs. 10-12. Assuming a suitable LADF is identified, at block 610 the query module 220 transmits the selected LADF (as LADF 134) to the requesting mobile device 104.

In response to receiving the LADF 134 in response to the LADF request 236, the 15 request module 360 provides the LADF 134 to both the localization module 362 and the feedback module 364. At block 612, the localization module 362 employs one or more well-known loop closure algorithms to localize the mobile device 104 to the coordinate frame represented in the LADF 134 based on comparisons of spatial features from the area detected by the spatial feature detection module 358 to the sparse point cloud of spatial features represented in the LADF 134, the result of 20 which is a determine pose 384 (FIG. 3) of the mobile device 104 in the identified coordinate frame. Further, in the event that the LADF was able to be georeferenced, the determined pose 384 likewise may be georeferenced through the appropriate translation. With the mobile device 104 localized and the pose 384 determined, at block 614 one or more components of the mobile device 104 may provide various 25 functionality that relies on accurate pose information, such as the implementation of AR or VR content, multiple-player games, navigation tools, and the like.

The generation of the LADF 134 relies on the detection and selection of spatial 30 features observed in the area by one or more mobile devices 104. However, as the area may have transient objects included in the observed spatial features, or as the configuration of the area may change over time, an LADF 134 may become less useful, or “stale.” Accordingly, in some embodiments, with user permission, the mobile devices 104 may provide feedback on the LADF, which may be used by the

ADF server 102 to refine, or “refresh” the LADF for subsequent use by other mobile devices. To this end, at block 616 the feedback module 364 may compare spatial features identified by the spatial feature detection module 358 in the area with the spatial features of the sparse point cloud represented by the LADF 134 and generate

5 LADF feedback 238 representing one or more of: spatial features of the LADF 134 that were also observed by the spatial feature detection module 358; spatial features of the LADF 134 that were not observed by the spatial feature detection module 358; and spatial features observed by the spatial feature detection module 358 that were not present in the LADF 134. The LADF feedback 238 is uploaded to the ADF server
10 102. At block 618, the spatial feature filter module 216 receives the LADF feedback 238, updates the feature scores of spatial features in the feature score datastore 208 accordingly, and then may update the LADF to include new spatial features or to remove previously-included spatial features based on the updated feature scores.

FIGs. 7-9 together illustrate an example of the process of refining an LADF 134

15 based on LADF feedback 238 in accordance with at least one embodiment. Initially, all spatial features represented in an LADF 134 are provided an initial feature score, which may be the same for all spatial features in the LADF 134, or which may be based on one or more properties of the spatial feature. For example, a spatial feature may be assigned an initial feature score based on the number of times it
20 appeared in the ADFs of the ADF cluster from which the LADF 134 was generated. For the following description, it is assumed that a higher feature score reflects a more reliable spatial feature, and conversely a lower feature score indicates a less reliable spatial feature. The LADF feedback 238 supplied by a mobile device 104 can include one or more entries, with each entry either confirming the observation of a spatial
25 feature present in the LADF 134 by the mobile device 104, denying the observation of a spatial feature present in the LADF 134 by the mobile device, or suggesting a new spatial feature observed by the mobile device 104 and not present in the LADF 134. The feature scores of spatial features thus may be adjusted based on the entries and the feedback type they represent.

30 The method 700 of FIG. 7 represents the processing implemented by the ADF server 102 for each entry of the LADF feedback 238. At block 702, the spatial feature filter module 216 accesses a selected entry of the LADF feedback 238 to determine the

type of feedback represented by the entry. In the event that the feedback entry confirms that the mobile device 104 observed a corresponding spatial feature of the LADF 134, at block 704 the spatial feature filter module 216 extracts one or more parameters pertaining to the identification of the spatial feature as represented in the feedback entry, such as a time stamp indicating a time of observation of the spatial feature by the mobile device 104, and records the extracted parameters in the feature score datastore 208. Further, at block 706, the spatial feature filter module 216 increases the feature score for the corresponding spatial feature to reflect its recent observation.

5 Returning to block 702, in the event that the feedback entry indicates that the mobile device 104 did not observe a spatial feature represented in the LADF 134, at block 708 the spatial feature filter module 216 extracts one or more parameters pertaining to the identification of the spatial feature as represented in the feedback entry, such as a time stamp indicating a time when the mobile device 104 was capturing imagery of the region in which the spatial feature should have been observed but was not, and records the extracted parameters in the feature score datastore 208. Further, at block 710, the spatial feature filter module 216 decreases the feature score for the corresponding spatial feature to reflect its missed observation.

10 Returning again to block 702, in the event that the feedback entry indicates that the mobile device 104 reliably observed a spatial feature not represented in the LADF 134, at block 712 the spatial feature filter module 216 extracts one or more parameters pertaining to the identification of the previously-unobserved spatial feature as represented in the feedback entry, such as a time stamp indicating a time when the mobile device 104 first or last observed the spatial feature, a frequency of 15 observation of the spatial feature (e.g., the percentage of times the spatial feature was observed when the mobile device 104 was oriented in the direction of the region containing the spatial feature), and the like, and records the extracted parameters in the feature score datastore 208. Further, at block 714, the spatial feature filter module 216 either creates a feature score for the corresponding spatial feature in the feature score datastore 208 and assigns an initial value to the feature score if the spatial feature is not already represented in the feature score datastore 208. 20 Otherwise, if a feature score is already present for the newly-observed spatial feature

(that is, the spatial feature was also observed by another mobile device 104 previously), the spatial feature filter module 216 increases the feature score for the spatial feature. Method 700 then returns to block 702 to repeat the process for the next entry in the LADF feedback 238.

5 FIG. 8 illustrates an example method 800 for evaluating a candidate spatial feature for inclusion in an LADF 134 in accordance with at least one embodiment. As described above, mobile devices 104 may provide LADF feedback 238 that identifies newly observed spatial features in an area that may be candidates for inclusion in the LADF 134 for the area, and the spatial feature filter module 216 may create a feature score the first time a spatial feature is detected, and increase the feature score each time the same spatial feature is observed by another mobile device 104. Thus, the feature score for each of these spatial features serves as an indicator of the viability of the candidacy of the spatial feature for inclusion in the LADF 134 (that is, how "reliable" the spatial feature is for purposes of localization). Accordingly, in response to a trigger condition, such as the aging of an LADF 134 exceeding a certain threshold, or in response to sufficient feedback from mobile devices 104 indicating that a revision to the LADF 134 would be warranted, at block 802 the spatial feature filter module 216 selects one of the candidate spatial features under consideration for inclusion in the LADF 134 from the feature score datastore 208 and at block 804 the spatial feature filter module 216 compares the feature score of the selected candidate spatial feature with a specified threshold, designated "THRESH_H". This threshold may be a fixed threshold, or may be based on one or more current conditions. For example, in one embodiment, the threshold THRESH_H may be set to the median current feature score for the spatial features currently included in the LADF 134.

10 15 20 25 30

In the event that the feature score of the candidate spatial feature does not exceed the threshold THRESH_H, the spatial feature filter module 216 ceases any further consideration of the candidate spatial feature (in the current round of evaluation), and the method 800 returns to block 802 for the selection of the next candidate spatial feature. Otherwise, if the feature score exceeds the threshold THRESH_H, at block 806 the spatial feature filter module 216 signals the localization generation module 218 to incorporate the selected candidate spatial feature into the LADF 134. In some embodiments, the incorporation of a new spatial feature into the LADF 134 may

require the elimination of another spatial feature from the LADF 134. In such cases, the spatial feature filter module 216 may select, for example, the spatial feature having the lowest feature score of all of the spatial features presently included in the LADF 134. After incorporating the candidate spatial feature into the LADF 134, the 5 method 800 returns to block 802 for the selection and evaluation of the next candidate spatial feature.

FIG. 9 illustrates an example method 900 for evaluating a spatial feature for removal from an LADF 134 in accordance with at least one embodiment. As described above, mobile devices 104 may provide LADF feedback 238 that confirms whether spatial

10 features of the LADF 134 were observed by the mobile device 104, and the spatial feature filter module 216 may adjust the feature scores of the spatial features accordingly. Thus, the feature score for each of these spatial features serves as an indicator of the current reliability of the spatial feature for localization purposes.

Accordingly, in response to a trigger condition, at block 902 the spatial feature filter 15 module 216 selects a spatial feature currently included in the LADF 134 and at block 904 the spatial feature filter module 216 compares the feature score of the selected spatial feature with a specified threshold, designated "THRESH_L". This threshold may be a fixed threshold, or may be based on one or more current conditions. For example, in one embodiment, the threshold THRESH_L may be set to the median 20 current feature score for the candidate spatial features not presently included in the LADF 134.

In the event that the feature score of the candidate spatial feature does exceed the threshold THRESH_L, the spatial feature filter module 216 ceases any further consideration of the selected spatial feature (in the current round of evaluation), and

25 the method 900 returns to block 902 for the selection of the next spatial feature in the LADF 134. Otherwise, if the feature score falls below the threshold THRESH_L, at block 906 the spatial feature filter module 216 signals the localization generation module 218 to remove the selected spatial feature from the LADF 134. In some embodiments, the removal of a spatial feature from the LADF 134 may require the 30 selection of another spatial feature to replace the removed spatial feature and thus serve as a trigger condition to initiate the candidate spatial feature evaluation process of method 800 of FIG. 8. After removing the selected spatial feature from the LADF

134, the method 900 returns to block 902 for the selection and evaluation of the next spatial feature in the LADF 134.

The LADFs maintained by the ADF server 102 contain representations of visual features and geometric information of an area of the world. This information is

5 potentially sensitive, and thus the ADF server 102 takes care to prevent unauthorized or unintended access to the LADF content stored in the LADF datastore 206 and the ADF content stored in the ADF datastore 204 (which may comprise the same datastore in some implementations). To this end, the ADF server 102 employs a privacy safeguard in the form of a two-layer indexing scheme for the LADF content of
10 the LADF datastore 206. FIGs. 10-12 illustrate an example configuration of the query module 220 and the LADF datastore 206 and operations thereof in accordance with this two-layer indexing scheme. The ADF datastore 204 may be similarly configured in the manner described below.

FIG. 10 illustrates an example two-layer query interface 1000 implemented by the

15 query module 220 and the LADF datastore 206. As depicted, each LADF (e.g., LADF 134) stored by the LADF datastore 206 is stored in a corresponding LADF entry 1002 of a LADF database 1004 implemented in the LADF datastore 206. Each LADF entry 1002 includes a plurality of fields, including a UID field 1006 to store a
20 UID associated with the LADF stored therein, a spatial feature field 1008 to store the set of spatial features presently represented by the LADF, and a location field 1010 to store the location indicator or other location data associated with the LADF.

Without adequate safeguards, an unauthorized party may be able to obtain the LADF 134 from the LADF datastore 206, and thus risk exposure of potentially sensitive information. The two-layer query interface 1000 may more fully secure the LADF

25 data of the LADF datastore 206 by requiring that the requesting mobile device 104 or other requestor prove that they are at (or have been at) the area associated with a requested LADF. Some location indicators, such as GPS coordinates, may be easily spoofed or otherwise forged. Further, other types of location indicators, such as WLAN MAC addresses or cell tower BSIDs may cover regions far larger than the
30 area of interest, and thus may lack sufficient granularity. As such, location indicators alone may not be sufficient proof that a requesting mobile device is in fact in the specified area. Thus, the query interface 1000 may instead require proof that the

mobile device 104 is in the specified area by requiring that the mobile device 104 submit a set of spatial features observed by the mobile device 104 in the specified area. The query interface 1000 then may compare this set of spatial features with the spatial features of the LADF entries to identify a matching LADF. However, many 5 areas have similar structural configurations (e.g., similar rooms within an office building), and thus there is a risk that multiple LADFs may adequately match the submitted set of spatial features, and thus the incorrect LADF may be supplied to the requesting mobile device, which is doubly problematic as it both represents a potential privacy breach as well as provides the mobile device 104 with the incorrect 10 reference for localization.

Thus, to ensure both that the mobile device is in fact in the area it represents it is and to ensure the correct LADF is supplied to the mobile device, the query interface 1000 utilizes both a set of spatial features and one or more location indicators submitted by a mobile device to identify the correct LADF from the LADF database 1004. It should 15 be appreciated that under this approach, the mobile device or other requestor proves that the requestor is in the identified area by submitting spatial features observed in the identified area, and thus the release of the LADF for the identified area has a low risk of unintended disclosure of potentially sensitive information as the requestor simply may look around the indicated area to obtain visual information many orders of 20 magnitude more detailed than the visual information represented by the supplied LADF.

To this end, the query module 220 includes a spatial feature query interface 1012 and a location selection module 1014. The spatial feature query interface 1012 performs the initial query of candidate LADFs based on a search of the spatial 25 feature fields 1008 of the LADF entries 1002 of the LADF database 1004 using a spatial feature set 1016 in a LADF request 236 submitted by a mobile device 104. This search takes into account the relative geometry of the spatial features as they are projected into the image represented by LADF request 236 to check for consistency against those in the LADF database 1004. The location selection 30 module 1014 then selects from the identified candidate LADFs based on comparisons of the one or more location indicator 1018 submitted with the LADF

request 236 with the location information for each LADF in the location fields 1010 of the candidate LADFs.

Further, to prevent search or other access of the LADF entries 1002 using a location indicator or other key besides a set of spatial features, in some embodiments the

- 5 LADF datastore 206 is configured to index the LADF entries 1002 by the spatial feature fields 1008 and further is configured to avoid indexing of the LADF entries 1002 by the location fields 1010. In this configuration, the LADF entries 1002 of the LADF database 1004 are searchable via the spatial feature sets of the LADFs but cannot be initially searched by location. In this manner, submission of a location
- 10 indicator alone will not result in the identification or production of any LADFs 134 from the LADF database 1004.

FIG. 11 illustrates an example method 1100 of operation of the two-layer query interface 1000 of FIG. 10 in accordance with at least one embodiment of the present disclosure. The method 1100 initiates at block 1102 with the submission of an LADF request 236 for an indicated area from a mobile device or other requestor. In at least one embodiment, in order to accept a LADF request 236 for processing, the ADF

- 15 server 102 requires that the LADF request 236 include at least the spatial feature set 1016 observed in the indicated area and one or more location indicators 1018 obtained or determined by the mobile device in the indicated area.

- 20 At block 1104 the spatial feature query interface 1012 searches the LADF entries 1002 of the LADF database 1004 to find LADF entries 1002 that have spatial feature sets in the spatial feature fields 1008 that sufficiently overlap with the spatial feature set 1016 of the LADF request 236, and which have relative geometries that are consistent with the geometries of the projected spatial features of the spatial feature
- 25 set 1016. In one embodiment, for each LADF entry 1002 analyzed the spatial feature query interface 1012 may determine a count of matching spatial features between the spatial feature set 1016 and the spatial feature field 1008 of the LADF entry 1002. The spatial feature query interface 1012 then may select N candidate LADFs (e.g., candidate LADFs 1021, 1022, 1023 of FIG. 10) based on the count values
- 30 determined for the LADF entries 1002. For example, in some instances every LADF having a count greater than a specified threshold may be selected as a candidate

LADF. In other instances, the spatial feature query interface 1012 may select the N LADFs with the highest count values as the N candidate LADFs.

At block 1106 the spatial feature query interface 1012 verifies that at least one candidate LADF has been identified from the LADF database 1004. If not, at block

5 1108 the query module 220 signals to the requesting mobile device that an LADF is unavailable for the indicated area. Otherwise, the one or more identified candidate LADFs are supplied to the location selection module 1014 (e.g., by identifying the UIDs of the candidate LADFs), and at block 1110 the location selection module 1014 compares the one or more location indicators 1018 of the LADF request 236 with the 10 locations indicated by the location fields 1010 of the candidate LADFs to score a match between each candidate LADF and the one or more location indicators 1018.

At block 1112, the location selection module 1014 verifies whether there is a sufficient match between the one or more location indicators 1018 and the location information of one of the candidate LADFs. If not, at block 1114 the query module

15 220 signals to the requesting mobile device that an LADF is unavailable for the indicated area. Otherwise, at block 1116 the location selection module 1014 selects the candidate LADF with the best location match score as the LADF co-located as the identified area and provides this selected LADF as the LADF 134 for the area to the mobile device in response to the LADF request 236.

20 In addition to providing protection against unauthorized or unintended access to the LADFs of the LADF database 1004, performing the LADF query using spatial features and their geometry as detected by the mobile device ensures that the LADF ultimately selected through the two-stage query process will allow the mobile device to successfully localize using the matching LADF. This is because the first stage of 25 the LADF query process performs the same or similar “localization” process to identify candidate LADFs as the localization process performed at the mobile to find the pose of the mobile device relative to the LADF using the observed spatial features and their image geometry in the LADF.

30 FIG. 12 illustrates an alternative configuration of the two-layer query interface 1000 in accordance with at least one embodiment. As with the implementation depicted in FIG. 10, the LADF datastore 206 implements an LADF database 1204 comprising a plurality of LADF entries 1202, each LADF entry 1202 having the UID field 1006, the

spatial feature field 1008, and the location field 1010. However, in contrast to the implementation of FIG. 10, the LADF database 1204 instead indexes the LADF entries 1202 based on the location field 1010. Further, in this embodiment, the query module 220 comprises a location query interface 1212 to the LADF database 1204 and a spatial feature selection module 1214.

5 In the depicted implementation, the two-layer query for an LADF is implemented as a first layer search of LADF entries 1202 by the location query interface 1212 for LADFs that have location information that sufficiently matches the one or more location indicators 1018 of the LADF request 236 to identify a set of one or more 10 candidate LADFs. The spatial feature selection module 1214 then compares the spatial feature set 1016 of the LADF request 236 to select the candidate LADF that best matches the spatial feature set 1016 as the LADF 134 to be supplied to the requesting mobile device in response to the LADF request 236.

15 Much of the inventive functionality and many of the inventive principles described above are well suited for implementation with or in integrated circuits (ICs) such as application specific ICs (ASICs). It is expected that one of ordinary skill, notwithstanding possibly significant effort and many design choices motivated by, for example, available time, current technology, and economic considerations, when 20 guided by the concepts and principles disclosed herein will be readily capable of generating such ICs with minimal experimentation. Therefore, in the interest of brevity and minimization of any risk of obscuring the principles and concepts according to the present disclosure, further discussion of such software and ICs, if any, will be limited to the essentials with respect to the principles and concepts within the preferred embodiments.

25 In this document, relational terms such as first and second, and the like, may be used solely to distinguish one entity or action from another entity or action without necessarily requiring or implying any actual such relationship or order between such entities or actions. The terms "comprises," "comprising," or any other variation thereof, are intended to cover a non-exclusive inclusion, such that a process, 30 method, article, or apparatus that comprises a list of elements does not include only those elements but may include other elements not expressly listed or inherent to such process, method, article, or apparatus. An element preceded by "comprises ...

a" does not, without more constraints, preclude the existence of additional identical elements in the process, method, article, or apparatus that comprises the element.

The term "another", as used herein, is defined as at least a second or more. The terms "including" and/or "having", as used herein, are defined as comprising. The

5 term "coupled", as used herein with reference to electro-optical technology, is defined as connected, although not necessarily directly, and not necessarily mechanically.

The term "program", as used herein, is defined as a sequence of instructions designed for execution on a computer system. A "program", or "computer program", may include a subroutine, a function, a procedure, an object method, an object

10 implementation, an executable application, an applet, a servlet, a source code, an object code, a shared library/dynamic load library and/or other sequence of instructions designed for execution on a computer system.

The specification and drawings should be considered as examples only, and the scope of the disclosure is accordingly intended to be limited only by the following

15 claims and equivalents thereof. Note that not all of the activities or elements described above in the general description are required, that a portion of a specific activity or device may not be required, and that one or more further activities may be performed, or elements included, in addition to those described. Still further, the order in which activities are listed are not necessarily the order in which they are

20 performed. The steps of the flowcharts depicted above can be in any order unless specified otherwise, and steps may be eliminated, repeated, and/or added, depending on the implementation. Also, the concepts have been described with reference to specific embodiments. However, one of ordinary skill in the art appreciates that various modifications and changes can be made without departing

25 from the scope of the present disclosure as set forth in the claims below.

Accordingly, the specification and figures are to be regarded in an illustrative rather than a restrictive sense, and all such modifications are intended to be included within the scope of the present disclosure.

Benefits, other advantages, and solutions to problems have been described above

30 with regard to specific embodiments. However, the benefits, advantages, solutions to problems, and any feature(s) that may cause any benefit, advantage, or solution to

occur or become more pronounced are not to be construed as a critical, required, or essential feature of any or all the claims.

What is claimed is:

1. A method comprising:

storing, at a datastore of a computing system, a plurality of localization area description files, each localization area description file representing a corresponding set of spatial features and location information identifying an associated with the set of spatial features;

5 receiving, at the computing system, a request for a localization area description file from a mobile device, the request comprising a set of spatial features and at least one non-image location indicator;

10 identifying a set of one or more candidate localization area description files from the datastore based on one of the set of spatial features of the request; and the at least one location indicator of the request; and

15 selecting a localization area description file from the set of one or more candidate localization area description files based on the other of the set of spatial features of the request; and the at least one location indicator of the request; and

providing the selected localization area description file to the mobile device in response to the request.

2. The method of claim 1, wherein the at least one location indicator comprises an indicator of a geographical location of the mobile device.

20 3. The method of claim 2, wherein the at least one indicator comprises an identifier associated with a wireless base station.

4. The method of claim 1, wherein:

identifying the set of one or more candidate localization area description files comprises identifying the set of one or more candidate localization area description 25 files based on the set of spatial features of the request; and

selecting the localization area description file comprises selecting the localization area description file based on the at least one location indicator of the request.

5. The method of claim 4, wherein identifying the set of one or more candidate localization area description files based on the set of spatial features of the request comprises:

determining, for each localization area description file of the plurality of localization

5 area description, a first score based on a number of spatial features of the

localization area description file that match spatial features of the set of spatial features of the request; and

selecting the set of one or more candidate localization area description files based on the first score determined for each localization area description file of the plurality of

10 localization area description files.

6. The method of claim 5, wherein determining the first score for a localization area description file comprises:

determining the first score using a loop closure algorithm.

7. The method of claim 5, wherein selecting the localization area description file

15 based on the at least one location indicator comprises:

determining, for each candidate localization area description file, a second score based on a number of location indicators represented by the location information of the localization area description file that correspond to a location indicator of the one or more location indicators of the request; and

20 selecting the localization area description file based on the second score determined for each candidate localization area description file.

8. The method of claim 1, wherein:

identifying the set of one or more candidate localization area description files

comprises identifying the set of one or more candidate localization area description

25 files based on the one or more location indicators of the request; and

selecting the localization area description file comprises selecting the localization area description file based on the set of spatial features of the request.

9. The method of claim 8, wherein identifying the set of one or more candidate localization area description files based on the one or more location indicators of the request comprises:

determining, for each localization area description file of the plurality of localization

5 area description files, a first score based on a number of location indicators represented by the location information of the localization area description file that correspond to a location indicator of the one or more location indicators of the request; and

selecting the set of one or more candidate localization area description files based on

10 the first score determined for each localization area description file of the plurality of localization area description files.

10. The method of claim 9, wherein selecting the localization area description file based on the set of spatial features of the request comprises:

determining, for each candidate localization area description file, a second score

15 based on a number of spatial features of the candidate localization area description file that match spatial features of the set of spatial features of the request; and

selecting the localization area description file based on the second score determined for each candidate localization area description file.

11. A computing system comprising:

20 a datastore to store a plurality of localization area description files, each localization area description file representing a corresponding set of spatial features and location information identifying an associated with the set of spatial features;

a network interface to receive a request for a localization area description file from a mobile device, the request comprising a set of spatial features and at least one non-

25 image location indicator; and

a query module coupled to the datastore and the network interface, the query module comprising:

a query interface to identify a set of one or more candidate localization area description files from the datastore based on one of the set of spatial features of the request; and the at least one location indicator of the request; and

a selection module to select a localization area description file from the set of one or more candidate localization area description files based on the other of the set of spatial features of the request; and the at least one location indicator of the request; and

wherein the query module is to provide the selected localization area description file to the mobile device via the network interface in response to the request.

10 12. The computing system of claim 11, wherein the at least one location indicator comprises at least one of: an indicator of a geographical location of the mobile device; and an identifier associated with a wireless base station.

13. The computing system of claim 11, wherein:

the query interface is to identify the set of one or more candidate localization area

15 description files based on the set of spatial features of the request; and

the selection module is to select the localization area description file based on the at least one location indicator of the request.

14. The computing system of claim 13, wherein:

the datastore is to index each localization area description file based on the set of

20 spatial features of the localization area description file without indexing the localization area description file based on the location information of the localization area description file.

15. The computing system of claim 13, wherein:

the query interface is to identify the set of one or more candidate localization area

25 description files by:

determining, for each localization area description file of the plurality of localization area description, a first score based on a number of spatial features of the

localization area description file that match spatial features of the set of spatial features of the request; and

selecting the set of one or more candidate localization area description files based on the first score determined for each localization area description file of the plurality of

5 localization area description files.

16. The computing system of claim 15, wherein:

the selection module is to determine the first score for a localization area description file using a loop closure algorithm.

17. The computing system of claim 15, wherein:

10 the selection module is to select the localization area description file by:

determining, for each candidate localization area description file, a second score based on a number of location indicators represented by the location information of the localization area description file that correspond to a location indicator of the one or more location indicators of the request; and

15 selecting the localization area description file based on the second score determined for each candidate localization area description file.

18. The computing system of claim 11, wherein:

the query interface is to identify the set of one or more candidate localization area description files based on the one or more location indicators of the request; and

20 the selection module is to select the localization area description file based on the set of spatial features of the request.

19. The computing system of claim 18, wherein:

the query interface is to identify the set of one or more candidate localization area description files by:

25 determining, for each localization area description file of the plurality of localization area description, a first score based on a number of location indicators represented

by the location information of the localization area description file that correspond to a location indicator of the one or more location indicators of the request; and

selecting the set of one or more candidate localization area description files based on the first score determined for each localization area description file of the plurality of

5 localization area description files.

20. The computing system of claim 19, wherein:

the selection module is to select the localization area description file by:

determining, for each candidate localization area description file, a second score based on a number of spatial features of the candidate localization area description

10 file that match spatial features of the set of spatial features of the request; and

selecting the localization area description file based on the second score determined for each candidate localization area description file.

21. A mobile device comprising:

at least one imaging sensor to capture imagery of an area in which the mobile device

15 is located;

a spatial feature detection module coupled to the at least one imaging sensor, the spatial feature detection module to determine a set of spatial features in the captured imagery;

a request module coupled to the spatial feature detection module, the request


20 module to generate a request for a localization area description file for the area, the localization area description file including the set of spatial features and one or more location indicators associated with the area; and

a network interface coupled to the request module, the network interface to transmit the request to a remote computing system.

25 22. The mobile device of claim 21, further comprising:

a localization module coupled to the network interface, the spatial feature detection module, and the request module, the localization module to perform a localization process using a spatial feature detected by the spatial feature detection module and a localization area description file received from the remote computing system in

5 response to the request.

**FIG. 1**





FIG. 2

3/11







**FIG. 4**

FIG. 5



6/11

**FIG. 6**


7/11


**FIG. 7**

8/11

**FIG. 8****FIG. 9**


9/11



**FIG. 11**

11/11

FIG. 12



INTERNATIONAL SEARCH REPORT

International application No
PCT/US2016/030956

A. CLASSIFICATION OF SUBJECT MATTER
INV. H04W64/00 H04W4/02
ADD. G01S5/02 G06T7/00

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)
H04W G01S G06T G01C G06K G06F

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)

EPO-Internal, WPI Data

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 2014/323148 A1 (SCHMALSTIEG DIETER [AT] ET AL) 30 October 2014 (2014-10-30) paragraphs [0002] - [0012], [0021] - [0081]; figures 1-6 ----- US 2014/357290 A1 (GRABNER MICHAEL [US] ET AL) 4 December 2014 (2014-12-04) paragraphs [0001], [0003], [0014] - [0058]; figures 1A-9 ----- -/-	1-22 1-22

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
"E" earlier application or patent but published on or after the international filing date
"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)
"O" document referring to an oral disclosure, use, exhibition or other means
"P" document published prior to the international filing date but later than the priority date claimed

"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

"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

"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

"&" document member of the same patent family

Date of the actual completion of the international search	Date of mailing of the international search report
21 July 2016	28/07/2016
Name and mailing address of the ISA/ European Patent Office, P.B. 5818 Patentlaan 2 NL - 2280 HV Rijswijk Tel. (+31-70) 340-2040, Fax: (+31-70) 340-3016	Authorized officer Jaster, Nicole

INTERNATIONAL SEARCH REPORT

International application No
PCT/US2016/030956

C(Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 2014/254934 A1 (LAXMINARAYANA BHAT PRASAD HORABAILU [IN] ET AL) 11 September 2014 (2014-09-11) paragraphs [0003], [0005] - [0007], [0009] - [0013], [0015] - [0037], [0049] - [0114]; figures 1-7B -----	1-22
X	US 2013/222369 A1 (HUSTON CHARLES D [US] ET AL) 29 August 2013 (2013-08-29) paragraphs [0003] - [0013], [0036] - [0075], [0079] - [0089], [0092] - [0197]; figures 1A-9, 11A-11D, 13A-18 -----	21,22
A	US 2012/315884 A1 (FORUTANPOUR BABAK [US] ET AL) 13 December 2012 (2012-12-13) paragraphs [0001], [0002], [0004] - [0018], [0032] - [0059]; figures 1-12 -----	1-20
A		1-22

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No

PCT/US2016/030956

Patent document cited in search report	Publication date	Patent family member(s)			Publication date
US 2014323148	A1 30-10-2014	CN 105143821 A EP 2992299 A1 KR 20160003731 A US 2014323148 A1 WO 2014179297 A1			09-12-2015 09-03-2016 11-01-2016 30-10-2014 06-11-2014
US 2014357290	A1 04-12-2014	AU 2014274343 A1 CA 2910355 A1 CN 105408762 A EP 3004915 A1 KR 20160016808 A US 2014357290 A1 WO 2014193873 A1			12-11-2015 04-12-2014 16-03-2016 13-04-2016 15-02-2016 04-12-2014 04-12-2014
US 2014254934	A1 11-09-2014	NONE			
US 2013222369	A1 29-08-2013	CA 2864003 A1 CN 104641399 A EP 2817785 A2 KR 20140127345 A US 2013222369 A1 US 2015287241 A1 US 2015287246 A1 WO 2013126784 A2			29-08-2013 20-05-2015 31-12-2014 03-11-2014 29-08-2013 08-10-2015 08-10-2015 29-08-2013
US 2012315884	A1 13-12-2012	CN 103703799 A EP 2719206 A1 JP 2014524168 A JP 2016076253 A KR 20140031955 A US 2012315884 A1 WO 2012170739 A1			02-04-2014 16-04-2014 18-09-2014 12-05-2016 13-03-2014 13-12-2012 13-12-2012