



- (51) International Patent Classification:
G06K 9/46 (2006.01)
- (21) International Application Number:
PCT/US2014/033401
- (22) International Filing Date:
8 April 2014 (08.04.2014)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:
13/858,813 8 April 2013 (08.04.2013) US
- (71) Applicant: **AMAZON TECHNOLOGIES, INC.**
[US/US]; P.O. Box 8102, Reno, Nevada 89507 (US).
- (72) Inventors: **RAMASWAMY, Sharadh**; 410 Terry Ave. North, Seattle, Washington 98109 (US). **BELL, Matthew Paul**; 410 Terry Ave. North, Seattle, Washington 98109 (US). **KHOKLOV, Dmitri**; 410 Terry Ave. North, Seattle, Washington 98109 (US). **STAFFORD, David W.**; 410 Terry Ave. North, Seattle, Washington 98109 (US). **NOBLE, Isaac Scott**; 410 Terry Ave. North, Seattle, Washington 98109 (US).
- (74) Agents: **LOHR, Jason D.** et al.; Novak Druce Connolly Bove + Quigg LLP, 1000 Louisiana Street, Fifty-Third Floor, Houston, Texas 77002 (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, LT, 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, 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).

[Continued on next page]

(54) Title: AUTOMATIC RECTIFICATION OF STEREO IMAGING CAMERAS




FIG. 4(e)

(57) Abstract: An electronic device can have two or more pairs of cameras capable of performing three-dimensional imaging. In order to provide accurate disparity information, these cameras should be sufficiently rectified. Automatic rectification can be performed by periodically capturing images with the cameras of interest, and locating matching feature points in corresponding images captured by those cameras. Small misalignment errors can be treated as linear translations, such that a set of linear equations can be used to solve for the misalignments. Another process can process a set of homographies for the cameras until a cost function converges. Various other approaches can be used as well, such as to directly solve for yaw, pitch, and roll errors. Once this information is obtained, the misalignment values (or related values) can be stored for use in correcting images subsequently captured by those cameras.

**Declarations under Rule 4.17:**

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

Published:

- *with international search report (Art. 21(3))*
- *before the expiration of the time limit for amending the claims and to be republished in the event of receipt of amendments (Rule 48.2(h))*

AUTOMATIC RECTIFICATION OF STEREO IMAGING CAMERAS

BACKGROUND

[0001] People are increasingly interacting with computers and other electronic devices in new and interesting ways. For example, mobile devices are increasingly offering multiple high quality cameras that enable additional types of functionality. In some devices, one or more pairs of these high quality cameras can be used to provide three-dimensional (“3D”) image capture, such as stereoscopic image capture, for both still and video imaging. A pair of cameras capturing three-dimensional images is offset by an amount that causes objects represented in an image captured by each camera to be located in slightly different locations, as the objects are imaged from slightly different points of view. The difference in location between the images, commonly referred to as the amount of disparity, is what gives a three-dimensional image its apparent depth when displayed to a user, as the amount of disparity changes with distance of an object from the cameras. Based on information such as the amount of offset between the cameras, the disparity also can be used to provide a measure of distance to each of the objects represented in the three-dimensional image. In order to produce an accurate 3D image with accurate distance information, however, the cameras have to be properly aligned, such as to have substantially parallel optical axes. Unfortunately, the cameras can become out of alignment over time due to effects such as shock and mechanical wear. While calibration procedures exist to compensate for misalignments, these procedures typically require user interaction and a calibration object, which is not convenient in many situations, and does not allow for continued, incremental adjustments over time unless the user frequently recalibrates the camera system.

25

BRIEF DESCRIPTION OF THE DRAWINGS

[0002] Various embodiments in accordance with the present disclosure will be described with reference to the drawings, in which:

[0003] **FIG. 1** illustrates front and back views of an example computing device including a conventional pair of cameras for 3D imaging that can be utilized in accordance with various embodiments;

30

[0004] FIG. 2 illustrates front and back views of an example computing device including a multiple cameras capable of providing 3D imaging in accordance with various embodiments;

[0005] FIGS. 3(a), (b), (c), and 3(d) illustrate examples of images that can be used
5 and/or generated using stereoscopic cameras in accordance with various embodiments;

[0006] FIG. 4(a), 4(b), 4(c), 4(d), and 4(e) illustrate examples of locating feature points in an image and determining misalignment errors that can be utilized in accordance with various embodiments;

[0007] FIG. 5 illustrates a first example process for determining misalignment errors
10 that can be analyzed in accordance with various embodiments;

[0008] FIG. 6 illustrates a second example process for determining misalignment errors that can be analyzed in accordance with various embodiments;

[0009] FIG. 7 illustrates an example configuration of components of a computing device such as that illustrated in FIG. 2; and

15 [0010] FIG. 8 illustrates an example environment in which various embodiments can be implemented.

DETAILED DESCRIPTION

[0011] Systems and methods in accordance with various embodiments of the present disclosure may overcome one or more of the aforementioned and other deficiencies
20 experienced in conventional approaches to acquiring image information using an electronic device. In particular, various embodiments enable the capture of three-dimensional (3D) image data with accurate disparity information. In order to provide the accurate disparity information, various approaches can automatically correct for misalignments of any of a set of cameras. Such approaches can be performed unknown
25 to a user, in at least some embodiments, and do not require a known calibration object or manual intervention on the part of the user.

[0012] In one embodiment, two or more pairs of cameras can be rectified by capturing images using each of the cameras and locating feature points (or interest points, corners, etc.) that appear in each of those images. Image rectification generally refers to a
30 transformation process used to project two or more images onto a common image plane,

correcting for image distortion by transforming the image into a standard coordinate system. Rectifying cameras, then, refers to determining misalignments or other aspects of those cameras that can result in such image distortion, and using this information to correct for resulting distortions in images subsequently captured by those cameras. For
5 small errors, misalignments of any or all the cameras can be treated as linear translations in the images. Accordingly, a set of linear equations can be used to solve for the misalignments of each camera, as determined by the coordinates of the located feature points. In other embodiments, a set of homographies can be determined for the cameras using the coordinates of the feature points. By iterating over the homographies until a
10 cost function converges, the misalignments of the cameras can be solved together. Other approaches can be used as well, such as to attempt to directly solve for yaw, pitch, and roll errors. Once this information is obtained, the misalignment values (or related values) can be stored for use in correcting images subsequently captured by those cameras. The process can be repeated at relatively short intervals in order to ensure that
15 misalignments between intervals remain relatively small, and to ensure that the cameras are sufficiently rectified at all times.

[0013] Various other applications, processes, and uses are presented below with respect to the various embodiments.

[0014] As mentioned above, conventional electronic devices offer 3D imaging using a
20 pair of high resolution matched cameras. For example, **FIG. 1** illustrates front and back views of an example electronic device **100** of the prior art wherein a front of the device includes a display screen **102** and other input elements, such as a front-facing camera **104** for video chat or other such purposes. The display can be, for example, a glasses-free 3D capable display, a display capable of presenting glasses-assisted 3D image
25 information, or a conventional 2D display. The device also includes a pair of matched cameras **106**, **108** on the back side of the device. These cameras are separated a distance sufficient to enable three-dimensional imaging and typically are relatively high resolution cameras (*e.g.*, 5.0MP or above for conventional devices). Processes discussed herein can be used to rectify a pair of cameras, accounting for misalignments
30 along at least one axis.

[0015] **FIG. 2** illustrates front and back views of another electronic device **200** that can be utilized in accordance with various embodiments. Although a portable

computing device (*e.g.*, a smart phone, an e-book reader, or tablet computer) is shown, it should be understood that various other types of electronic devices that are capable of determining and processing input can be used in accordance with various embodiments discussed herein. These devices can include, for example, notebook computers,

5 personal data assistants, cellular phones, video gaming consoles or controllers, and portable media players, among others. The electronic device can utilize some of the same elements as a conventional device, such as may include a display screen **202** and at least one front-facing camera **204**. This example also includes four cameras **208**, **210**, **212**, **214** arranged to provide at least two stereoscopic imaging pairs. These

10 cameras are labeled C1, C2, C3, and C4 for reference, as used later herein. While shown on the backside of the device in this example, it should be understood that the pairs could be on the front of the device or in one or more other appropriate locations. Further, while two pairs of stereoscopic cameras are illustrated, it should be understood that there can be additional pairs of cameras utilized as well as discussed elsewhere

15 herein. The example device can include other elements useful for imaging as well, such as a light sensor **216** for determining an amount of ambient light and a white light LED **218**, or other such illumination element, useful in illuminating objects within at least a portion of a field of view of at least one of the cameras **208**, **210**, **212**, **214**. Each image capture element may be, for example, a camera, a complimentary metal-oxide

20 semiconductor (CMOS) device, or another appropriate image capturing element or sensor. It should be understood that while certain elements are shown to be included on a “front” or “back” side of the device that any or all of these elements can be positioned on various sides, edges, faces, or other regions of such a device. Further, terms such as “front,” “back,” and “top” are used for purposes of explanation and are not intended to

25 be interpreted as required orientations unless otherwise stated. Further still, while terms such as “rectangular grid” or “rectangular pattern” are used herein to describe the relative arrangements of various cameras, it should be understood that pairs of cameras in such a system are positioned along orthogonal axes, such as horizontal and vertical axes, such that a camera of a pair is positioned horizontally or vertically (or along other

30 orthogonal axes) with respect to another camera of the pair. It should be noted that the cameras do not have to form a proper rectangle, but can form other patterns such as a cross, set of parallel lines, points along a rectilinear grid, etc. Various other geometries and arrangements can be used as well within the scope of the various embodiments.

[0016] For any pair of these cameras that have at least a partially overlapping field of view, three-dimensional imaging can be performed by capturing image information for one or more objects from two different perspectives or points of view, and combining the information to produce a 3D image. In at least some embodiments, the fields of view can initially be matched through careful placement and calibration, such as by imaging a known calibration standards and adjusting an optical axis of one or more cameras to have those axes be substantially parallel. In other embodiments, the calibration process does not adjust the optical axis of any cameras, but instead provides an adjustment to be applied to any captured image data in order to account for and/or remove any detected misalignment. Methods for producing a 3D image using image information from different perspectives are well known in the art and will not be described in detail herein. Example approaches include calculating an amount of disparity through a process such as edge matching, feature location and matching, color matching, and/or texture matching, and combining the disparity information with color information from each perspective to generate a three-dimensional image, either before or at a time of display. For example, if the image information is matched then the image information can be combined and/or displayed directly on a 3D-capable display, wherein the human brain can effectively do at least some of the 3D processing. In other examples, the image information can be otherwise combined or processed at the time of display such that upon displaying the information a 3D image is generated. It should be understood, however, that 3D image data can be used for other purposes or for further processing, such that using the image data to generate and display a 3D image is not required. For example, the data can be used to determine shape and/or relative position information for various computer vision techniques.

[0017] For example, **FIG. 3(a)** illustrates what will be referred to herein as a “left” image **300** and **FIG. 3(b)** illustrates what will be referred to herein as a “right” image **320**, although other orientations can be used as well in other situations and embodiments. These images are each captured by one of a pair of cameras that are offset from each other along a horizontal axis, such that one camera is offset to the left and one offset to the right in this orientation. As discussed elsewhere herein, pairs of cameras can be offset in up and down directions, among other such options. Since each image is captured using a camera at a slightly different location, the position of objects in each image will be different. As an example, if one were to overlay one of these

images **300** on top of the other **320**, as illustrated in the example image **340** of **FIG. 3(c)**, it can be seen that each of the objects is slightly offset, with objects closer to the camera being more offset than others, due to disparity differences as discussed previously. When displayed and/or viewed as a three dimensional image, however, the
5 disparity causes the objects in the image to appear to be located at different depths, as illustrated in the image **360** of **FIG. 3(d)**, resulting from the disparity between the two images.

[0018] In order for the images to combine to form an accurate three-dimensional image, the cameras used to capture the component images should be sufficiently aligned
10 and/or rectified to represent the correct amount of disparity. Misalignments along the direction of the offset can cause objects to have an incorrect amount of disparity, which can affect the apparent depth or distance to the object. Misalignments in other directions can cause various potential problems, such as problems with processing computer vision algorithms, problems with objects being blurry or otherwise improperly
15 rendered when the component images are combined for the three-dimensional image, etc. Accordingly, it can be desired to account for, and remove, any effects of the misalignment of the cameras. As mentioned, cameras can be increasingly misaligned over time, due to factors such as impact or shock to the device, natural degradation (*e.g.*, due to temperature and/or humidity), vibration over time, mechanical wear or failure,
20 and other such causes.

[0019] As discussed, it can be desirable to enable a computing device or electronic device including the cameras to be able to automatically rectify the stereo cameras without manual intervention by the user, although manual intervention or triggering can be used in some embodiments. It also can be desirable to enable the rectifying to be
25 done without the need for a calibration object or other such item.

[0020] Accordingly, approaches in accordance with various embodiments can utilize images captured of random (or other) objects to attempt to rectify the stereo camera pairs. The rectifying can be performed at any appropriate time, such as at regular intervals, at random times, or in response to detected events, such as rapid movements
30 or force as detected by a motion sensor or other component of the device. Where there are multiple pairs of stereo cameras on a device, different combinations of the cameras can be used to improve the rectification of each camera.

[0021] In one example, images can be simultaneously and/or concurrently captured by at least one pair of the cameras. Feature points, image points, or other such features of each image can be determined using one or more feature detection algorithms, as may recognize unique features or shapes, among other such features, that can be recognized in each image. As an example, **FIG. 4(a)** illustrates an image **400** including a view of a hand **402** of a person. If this image is fed to a feature point recognition algorithm, or other such process, a unique feature point such the tip **404** of the user's thumb can be located in each image. As known for such purposes, the feature can be located due to the presence of a transition point, recognized pattern, abrupt change in color or intensity, or other such aspect of the image at that location. Once a particular feature point is located in each of the images, the relative position of that point in the images can be analyzed.

[0022] As an example, consider the image **420** represented in **FIG. 4(b)**, which represents an overlay of images captured by a pair of stereo cameras offset horizontally, such as cameras C1 and C2 in **FIG. 2**. As can be seen, there is an expected amount of offset or disparity **424** due to the separation of the cameras along the horizontal axis (or x-axis). If the cameras are properly aligned, however, there will be no appreciable offset along the vertical axis (or y-axis). When properly rendered or viewed, the component parts of **FIG. 4(b)** then can provide a proper three-dimensional image.

[0023] It may be the case, however, that there is some amount of misalignment of at least one of those cameras. For example, **FIG. 4(c)** illustrates an overlay **440** of images from the horizontally offset cameras as in **FIG. 4(b)**. In this example, however, there is additional offset **442** along the x-axis, which can affect the apparent distance to that object in the three-dimensional image. This movement along the x-axis can be thought of as rotation along the y-axis, or "yaw" as it is often referred. Similarly, **FIG. 4(d)** illustrates an overlay **460** of images from the horizontally offset cameras as in **FIG. 4(b)**, where there is an offset **442** along the y-axis, due to misalignment of at least one of the cameras. This offset along the y-axis can be attributed to rotation along the x-axis, or "pitch." As previously stated, for horizontally offset cameras such as cameras C1 and C2, there should be no appreciable offset along the y-axis. Further, **FIG. 4(e)** illustrates an overlay **480** of images from the horizontally offset cameras as in **FIG. 4(b)**, where there is an offset **482** along the x-axis and an offset **484** along the y-axis,

due to misalignment of at least one of the cameras. The combined offsets could be a translation in the x/y plane, and thus might be addresses as a combination of the previous two offsets. In this situation, however, the offset changes for each feature point in the image, as the misalignment results in a rotation of the image. This rotation
5 along the z-axis is often referred to as “roll.”

[0024] In order to determine how to properly rectify the cameras, then, it is not enough to simply look at the translations along the x-axis and y-axis between the images, but to look at the translations of the various feature points in order to also account for any amount of rotation among the cameras. Such an approach enables
10 adjusting for offsets in roll, pitch, and yaw.

[0025] One approach that can be utilized in accordance with various embodiments takes into account the fact that the cameras will generally become misaligned slowly over a period of time, such that if a regular rectification process is employed the cameras will move only a slight amount between processing times. The presence of
15 only small angles then can allow for a small angle approximation, where offsets can be analyzed and accounted for without having to consider differences in changes due to distance variation with respect to the change in angle. Further, such a process can take advantage of the fact that cameras such as C1, C2, C3, and C4 in **FIG. 2** are aligned such that for any given pair of the cameras with a horizontal or vertical offset, either the
20 x-coordinates or y-coordinates of the feature points in images captured by those cameras should be the same, depending on the orientation (as discussed with respect to **FIG. 4(d)**). For example, in **FIG. 2** feature points in an image captured by C1 should have the same y-coordinates as the corresponding feature points in an image captured by C2, and the same goes for images capture by C3 with respect to C4. Similarly, feature
25 points in an image captured by C1 should have the same x-coordinates as the corresponding feature points in an image captured by C3, and the same goes for images capture by C2 with respect to C4. In order to reliably perform a rectification of the cameras, approaches in accordance with various embodiments can attempt to concurrently solve for yaw, pitch, and roll of each of the cameras of interest.

30 [0026] **FIG. 5** illustrates one example process **500** for performing dynamic rectification that can be utilized in accordance with various embodiments. It should be understood that, for any process discussed herein, there can be additional, fewer, or

alternative steps performed in similar or alternative orders, or in parallel, within the scope of the various embodiments unless otherwise stated. In this example, images are acquired **502** using each camera to be rectified. As discussed, this can include two or more pairs of cameras arranged in a pattern, such as a rectangle or regular array, such that different pairs of the cameras can be used to perform three-dimensional imaging, and the pairs can be selected in at least two different directions, which in at least some embodiments are orthogonal directions. Also, as discussed elsewhere herein, in at least some embodiments all cameras of interest do not have to capture a respective image at substantially the same time, but at least pairs of cameras to be rectified should capture images at substantially the same time. Each captured image can be analyzed to determine **504** a set of feature points or other such aspects or portions of the image that can be located in the other images. As known for such purposes, feature points can be determined using algorithms such as feature detection algorithms (*e.g.*, SIFT or SURF), corner finding algorithms, pattern matching algorithms, contour detection algorithms, and the like. For this example, the process uses a feature detection algorithm to locate specific points in each image. Corresponding feature points can then be located **506** in each of the captured images, to the extent those feature points are represented in each of those images.

[0027] Thus, a feature point corresponding to an object represented in an image captured by four cameras will have four pairs of coordinates, with an (x,y) or other such pair of coordinates representing the location of that feature point in each image. As discussed elsewhere herein, the examples utilize a single feature point for purposes of explanation, but it should be understood that multiple feature points will generally be analyzed and/or otherwise utilized in various analyses within the scope of the various embodiments.

[0028] Once the feature points are detected, the coordinates of those feature points in each image can be determined **508**. For purposes of explanation, each feature point in an image captured by camera #1 (C1 in the example of **FIG. 2**) will have representative coordinates of (x_1, y_1) , each corresponding feature point in an image captured by camera #2 (C2 in the example of **FIG. 2**) will have representative coordinates of (x_2, y_2) , etc. In the case where all the cameras are aligned, the feature points would have at least some similar values, such that for pitch and roll $y_1 = y_2$ and $y_3 = y_4$, and for yaw and roll $x_1 =$

10/30

x_3 and $x_2 = x_4$. These values indicate that the stereo pairs (C1,C2), (C3,C4), (C1,C3), and (C2,C4) are rectified. If there is any misalignment, however, at least one of these rules will be violated. As an example, the following cost equation:

$$J = (y_1 - y_2)^2 + (y_3 - y_4)^2 + (x_1 - x_3)^2 + (x_2 - x_4)^2$$

5 can be summed over all feature points. In perfect conditions, $J = 0$. With noise, $J \approx T$, a small noise threshold. When misaligned, however, $J \gg T$. Thus, some small variation in alignment might be allowed, but there can be a minimum misalignment threshold specified such that when J exceeds that threshold, an attempt can be made to adjust for the misalignment to bring J back to on the order of T or less.

10 **[0029]** If misalignments are detected and corrected relatively frequently, as may vary by device and/or design, for example, it can reasonably be assumed that in most cases the misalignment errors in yaw, pitch, and/or roll over the intervening time period (*e.g.*, an hour or a day) will be relatively small, such as with an angle of less than five degrees or $\pi/36$ radians. For such small angles, a small angle approximation can be used

15 whereby small angle variations can be modeled as small linear displacements. As an example, for small angles:

$$\sin(\theta) \approx \theta$$

$$\cos(\theta) \approx 1$$

$$\tan(\theta) \approx \theta$$

20 And, where an in-plane rotation (*i.e.*, roll) by t degrees on a point (x_1, y_1) returns a point (x_1r, y_1r) :

$$\begin{bmatrix} x_1r \\ y_1r \end{bmatrix} = \begin{bmatrix} \cos(roll) & -\sin(roll) \\ \sin(roll) & \cos(roll) \end{bmatrix} \begin{bmatrix} x_1 \\ y_1 \end{bmatrix}$$

such that for a small roll value:

$$\begin{bmatrix} x_1r \\ y_1r \end{bmatrix} = \begin{bmatrix} 1 & -roll \\ roll & 1 \end{bmatrix} \begin{bmatrix} x_1 \\ y_1 \end{bmatrix}$$

25 Small amounts of pitch variation in a stereo pair then can be modeled as a uniform relative displacement ty in the y -axis direction (the relative pitch being analyzed as, for a single pair of cameras, only a relative pitch variation can be determined without a fixed frame of reference):

11/30

$$ty = \text{focalLength} * \tan(\text{relativePitch})$$

and small amounts of yaw variation in a stereo pair can be modeled as a uniform relative displacement tx in the x -axis direction:

$$tx = \text{focalLength} * \tan(\text{relativeYaw})$$

- 5 [0030] The focal length is a retrievable characteristic of the cameras, and is analyzed at least in part because longer focal lengths generally are more sensitive to misalignments. It can be desirable to determine the roll and x,y translation values such that the following are satisfied:

- points in images from C1 and C2 match in the y -component
 10 points in images from C1 and C3 match in the x -component
 points in images from C3 and C4 match in the y -component
 points in images from C2 and C4 match in the x -component

- [0031] In order to determine the appropriate values, a feature point detected in an image captured by one of the cameras can be designated as (x_i, y_i) for camera C_i , where i
 15 $= 1, 2, 3, 4$ (or more for additional cameras). It is desired to find the values for roll_1 , roll_2 , roll_3 , and roll_4 , as well as for displacements (tx_1, ty_1) , (tx_2, ty_2) , (tx_3, ty_3) , and (tx_4, ty_4) . To solve for these parameters for all feature points, a set of linear equations can be determined based on the information above. In this example, these linear equations are given by the following:

- 20 $x_1 * \text{roll}_1 + y_1 + ty_1 = x_2 * \text{roll}_2 + y_2 + ty_2$
 $x_1 - y_1 * \text{roll}_1 + tx_1 = x_3 - y_3 * \text{roll}_3 + tx_3$
 $x_3 * \text{roll}_3 + ty_3 + y_3 = x_4 * \text{roll}_4 + ty_4 + y_4$
 $x_2 - y_2 * \text{roll}_2 + tx_2 = x_4 - y_4 * \text{roll}_4 + tx_4$

- [0032] The coordinates from the images (*e.g.*, three or more feature point coordinates)
 25 can be inserted **510** in the set of four linear equations. These equations then can be solved **512**, such as by using a least squares matching, robust estimation, Gauss-Newton, QR decomposition, or Singular Value Decomposition process. The results (*e.g.*, calibration parameters) then can be stored **514** as misalignment corrections to be applied to images captured by each of the cameras. For example, the calibration

parameters can include lateral offset amounts for small angles. In other cases, the calibration parameters can include misalignments in roll, pitch, and/or yaw for any or all of the cameras analyzed, such that disparity can be corrected for subsequent images. Various other types of correction or misalignment data can be stored as well within the scope of the various embodiments. If a device has a mechanical mechanism for realigning the cameras, this information can alternatively be provided to the mechanism instead of being stored for future image adjustments. In embodiments where the misalignments are relative, the values stored can be normalized and/or adjusted such that the minimum average amount of adjustment is applied to each camera while still accounting for misalignment. As discussed, the process can be repeated at relatively regular and/or short intervals, or at other appropriate times, in order to ensure that the cameras are relatively aligned and that any misalignment errors are small. The ability to periodically detect and solve for small misalignments in yaw, pitch, and/or roll misalignment errors enables the calibration parameters to be updated and applied over time in order to maintain acceptable stereo rectification.

[0033] FIG. 6 illustrates another example process 600 for performing dynamic rectification that can be utilized in accordance with various embodiments. As with the prior example, images are acquired 602 using each camera to be rectified, and each captured image can be analyzed to determine 604 feature points or other such aspects or portions of the image that can be located in the other images. Corresponding feature points can then be located 606 in each of the captured images, to the extent those feature points are represented in each of those images. Once the feature points are detected, the coordinates of those feature points in each image can be determined 608.

[0034] As discussed above, the small misalignments can be modeled as linear transformations. One way to express a linear mapping that warps one image to another, while maintaining straightness of lines, is through use of a homography, H . A homography, as used in general for purposes such as computer vision, is an invertible transformation (usually expressed as a 3×3 matrix) from a projective space to itself that maps straight lines to straight lines. A correction process for misalignment of the four cameras can then initialize or otherwise select 610 an initial set of four homographies H_1 , H_2 , H_3 , and H_4 (such as one identity matrix for each respective camera, or an estimated matrix based on previous information), and can process these homographies

such that the coordinates of the feature points are aligned. To perform the estimations, (x_1r, y_1r) can be set as the rectified co-ordinates obtained from (x_1, y_1) using homography H_1 applied to images from camera C_1 , given:

$$\begin{bmatrix} x_1r \\ y_1r \end{bmatrix} = H_1 \begin{bmatrix} x_1 \\ y_1 \end{bmatrix}$$

- 5 Similarly (x_2r, y_2r) can be so defined and used for H_2 , (x_3r, y_3r) for H_3 , and (x_4r, y_4r) for H_4 . In certain conventional approaches, a fundamental matrix between pairs is determined before estimating the homographies. In this example, the process would instead estimate four fundamental matrices. Unfortunately, estimating the fundamental matrix is extremely unreliable. Additionally, these approaches cannot constrain the first
10 row of the homography matrix, which could result in sub-optimal solutions being estimated.

- [0035]** Thus, approaches in accordance with various embodiments can make certain assumptions as discussed above. For example, it can be assumed that $y_1r = y_2r$, $y_3r = y_4r$, $x_1r = x_3r$, and $x_2r = x_4r$, independent of the relationships between the original pairs of
15 co-ordinates. As discussed, there are four homographies to be estimated, and these homographies affect each other since the device utilizes a coupled stereo system. Accordingly, an iterative approach can be taken. In this example, one of the homographies is selected **612** for analysis. The other homographies are then held fixed **614**, assuming them to be optimal. An attempt is then made to match **616** the selected
20 homography to the fixed homographies. In order to perform the matching, a new cost Jr can be defined as:

$$Jr = (y_1r - y_2r)^2 + (y_3r - y_4r)^2 + (x_1r - x_3r)^2 + (x_2r - x_4r)^2,$$

- which can be summed over the feature points. The matching feature points (x_1, y_1) , (x_2, y_2) , (x_3, y_3) and (x_4, y_4) can be provided as inputs, and the initial homographies H_1 , H_2 ,
25 H_3 , and H_4 can be assumed to be identity. Accordingly, the algorithm can be processed and/or optimized until Jr is sufficiently minimized for each of the homographies, when the others are held fixed (since each processing results in a reduction of Jr , as discussed elsewhere herein). Once the matching homography has been determined, another determination can be made as to whether the cost function Jr has converged **618**. In
30 some embodiments, this determination can be made only after the set of homographies

has had another iteration, while in other embodiments the convergence can be analyzed after each individual iteration. If the cost function has not converged, the process can continue for another iteration. In each of the steps J_r is reduced, such that the cost function will eventually converge. If the cost function is determined to converge, according to one or more convergence criteria known or used for such purposes, the results (including types of results discussed elsewhere herein) can be stored **620** as misalignment corrections to be applied to images captured by each of the cameras.

[0036] Other approaches can be used as well within the scope of the various embodiments. For example, hypothesis testing can be used for different yaw, pitch, and roll solutions found using non-derivative base cost minimization approaches. These approaches can include, for example, Nelder-Mead, simplex search, or globally optimal solutions such as simulated annealing. Such an approach can be utilized since, for example, the intrinsic parameters of the camera can be obtained. Such an approach is similar to that discussed with respect to **FIG. 6**, except that this approach attempts to solve directly for yaw, pitch, and roll values, as a different way of expressing the transformation that has been applied.

[0037] For at least some of these and other such processes, it can be desirable to perform the misalignment analysis over a large set of feature points, for a significant number of frames or images, in order to attempt to average out any noise issues. Further, the timing of the image capture, at least for pairs of cameras being analyzed, should be as close as possible in order to avoid motion- or time-based effects, which can negatively impact the results. In at least some embodiments, it can be desirable to implement a global shutter to attempt to coordinate the timing of the capturing of the various images to be analyzed. In some embodiments, a motion sensor, such as an inertial sensor or accelerometer, can be used to determine whether the device is moving, such that alignment adjustments are not performed when the device is moving and results might be impacted by the motion. Various other criteria or optimizations can be utilized with approaches discussed herein as well within the scope of the various embodiments.

[0038] **FIG. 7** illustrates an example set of basic components of a computing device **700**, such as the device **200** described with respect to **FIG. 2**. In this example, the device includes at least one central processor **702** for executing instructions that can be

stored in at least one memory device or element **704**. As would be apparent to one of ordinary skill in the art, the device can include many types of memory, data storage or computer-readable storage media, such as a first data storage for program instructions for execution by the processor **702**, the same or separate storage can be used for images or data, a removable storage memory can be available for sharing information with other devices, etc. The device typically will include some type of display element **706**, such as a touch screen, electronic ink (e-ink), organic light emitting diode (OLED) or liquid crystal display (LCD), although devices such as portable media players might convey information via other means, such as through audio speakers. In at least some embodiments, the display screen provides for touch or swipe-based input using, for example, capacitive or resistive touch technology.

[0039] As discussed, the device in many embodiments will include at least one image capture element **708**, such as one or more cameras that are able to image a user, people, or objects in the vicinity of the device. An image capture element can include, or be based at least in part upon any appropriate technology, such as a CCD or CMOS image capture element having a determined resolution, focal range, viewable area, and capture rate. The device can also include at least one dedicated gesture component, such as an IR sensor or detector, operable to capture information for use in determining gestures or motions of the user, which will enable the user to provide input through the portable device without having to actually contact and/or move the portable device. The device can include at least one motion component **710**, such as an accelerometer or inertial sensor to measure translation, an electronic compass to calibrate direction in up to three dimensions, an electronic gyroscope to determine and/or track changes in orientation, and other such elements. Motion determining elements can help in determining movement and using the determined movement to correct image information. The device also can include at least one illumination element **712**, as may include one or more light sources (*e.g.*, white light LEDs, IR emitters, or flashlamps) for providing illumination and/or one or more light sensors or detectors for detecting ambient light or intensity, etc.

[0040] The example device can include at least one additional input device able to receive conventional input from a user. This conventional input can include, for example, a push button, touch pad, touch screen, wheel, joystick, keyboard, mouse,

trackball, keypad or any other such device or element whereby a user can input a command to the device. These I/O devices could even be connected by a wireless infrared or Bluetooth or other link as well in some embodiments. In some embodiments, however, such a device might not include any buttons at all and might be controlled only through a combination of visual (*e.g.*, gesture) and audio (*e.g.*, spoken) commands such that a user can control the device without having to be in contact with the device.

[0041] As discussed, different approaches can be implemented in various environments in accordance with the described embodiments. For example, **FIG. 8** illustrates an example of an environment **800** for implementing aspects in accordance with various embodiments. As will be appreciated, although a Web-based environment is used for purposes of explanation, different environments may be used, as appropriate, to implement various embodiments. The system includes an electronic client device **802**, which can include any appropriate device operable to send and receive requests, messages or information over an appropriate network **804** and convey information back to a user of the device. Examples of such client devices include personal computers, cell phones, handheld messaging devices, laptop computers, set-top boxes, personal data assistants, electronic book readers and the like. The network can include any appropriate network, including an intranet, the Internet, a cellular network, a local area network or any other such network or combination thereof. Components used for such a system can depend at least in part upon the type of network and/or environment selected. Protocols and components for communicating via such a network are well known and will not be discussed herein in detail. Communication over the network can be enabled via wired or wireless connections and combinations thereof. In this example, the network includes the Internet, as the environment includes a Web server **806** for receiving requests and serving content in response thereto, although for other networks, an alternative device serving a similar purpose could be used, as would be apparent to one of ordinary skill in the art.

[0042] The illustrative environment includes at least one application server **808** and a data store **810**. It should be understood that there can be several application servers, layers or other elements, processes or components, which may be chained or otherwise configured, which can interact to perform tasks such as obtaining data from an

appropriate data store. As used herein, the term "data store" refers to any device or combination of devices capable of storing, accessing and retrieving data, which may include any combination and number of data servers, databases, data storage devices and data storage media, in any standard, distributed or clustered environment. The application server **808** can include any appropriate hardware and software for integrating with the data store **810** as needed to execute aspects of one or more applications for the client device and handling a majority of the data access and business logic for an application. The application server provides access control services in cooperation with the data store and is able to generate content such as text, graphics, audio and/or video to be transferred to the user, which may be served to the user by the Web server **806** in the form of HTML, XML or another appropriate structured language in this example. The handling of all requests and responses, as well as the delivery of content between the client device **802** and the application server **808**, can be handled by the Web server **806**. It should be understood that the Web and application servers are not required and are merely example components, as structured code discussed herein can be executed on any appropriate device or host machine as discussed elsewhere herein.

[0043] The data store **810** can include several separate data tables, databases or other data storage mechanisms and media for storing data relating to a particular aspect. For example, the data store illustrated includes mechanisms for storing content (e.g., production data) **812** and user information **816**, which can be used to serve content for the production side. The data store is also shown to include a mechanism for storing log or session data **814**. It should be understood that there can be many other aspects that may need to be stored in the data store, such as page image information and access rights information, which can be stored in any of the above listed mechanisms as appropriate or in additional mechanisms in the data store **810**. The data store **810** is operable, through logic associated therewith, to receive instructions from the application server **808** and obtain, update or otherwise process data in response thereto. In one example, a user might submit a search request for a certain type of item. In this case, the data store might access the user information to verify the identity of the user and can access the catalog detail information to obtain information about items of that type. The information can then be returned to the user, such as in a results listing on a Web page

that the user is able to view via a browser on the user device **802**. Information for a particular item of interest can be viewed in a dedicated page or window of the browser.

[0044] Each server typically will include an operating system that provides executable program instructions for the general administration and operation of that server and

5 typically will include computer-readable medium storing instructions that, when executed by a processor of the server, allow the server to perform its intended functions. Suitable implementations for the operating system and general functionality of the servers are known or commercially available and are readily implemented by persons having ordinary skill in the art, particularly in light of the disclosure herein.

10 [0045] The environment in one embodiment is a distributed computing environment utilizing several computer systems and components that are interconnected via communication links, using one or more computer networks or direct connections. However, it will be appreciated by those of ordinary skill in the art that such a system could operate equally well in a system having fewer or a greater number of components
15 than are illustrated in **FIG. 8**. Thus, the depiction of the system **800** in **FIG. 8** should be taken as being illustrative in nature and not limiting to the scope of the disclosure.

[0046] The various embodiments can be further implemented in a wide variety of operating environments, which in some cases can include one or more user computers or computing devices which can be used to operate any of a number of applications.

20 User or client devices can include any of a number of general purpose personal computers, such as desktop or laptop computers running a standard operating system, as well as cellular, wireless and handheld devices running mobile software and capable of supporting a number of networking and messaging protocols. Such a system can also include a number of workstations running any of a variety of commercially-available
25 operating systems and other known applications for purposes such as development and database management. These devices can also include other electronic devices, such as dummy terminals, thin-clients, gaming systems and other devices capable of communicating via a network.

30 [0047] Most embodiments utilize at least one network that would be familiar to those skilled in the art for supporting communications using any of a variety of commercially-available protocols, such as TCP/IP, FTP, UPnP, NFS, and CIFS. The network can be, for example, a local area network, a wide-area network, a virtual private network, the

Internet, an intranet, an extranet, a public switched telephone network, an infrared network, a wireless network and any combination thereof.

[0048] In embodiments utilizing a Web server, the Web server can run any of a variety of server or mid-tier applications, including HTTP servers, FTP servers, CGI
5 servers, data servers, Java servers and business application servers. The server(s) may also be capable of executing programs or scripts in response requests from user devices, such as by executing one or more Web applications that may be implemented as one or more scripts or programs written in any programming language, such as Java[®], C, C# or C++ or any scripting language, such as Perl, Python or TCL, as well as combinations
10 thereof. The server(s) may also include database servers, including without limitation those commercially available from Oracle[®], Microsoft[®], Sybase[®] and IBM[®].

[0049] The environment can include a variety of data stores and other memory and storage media as discussed above. These can reside in a variety of locations, such as on a storage medium local to (and/or resident in) one or more of the computers or remote
15 from any or all of the computers across the network. In a particular set of embodiments, the information may reside in a storage-area network (SAN) familiar to those skilled in the art. Similarly, any necessary files for performing the functions attributed to the computers, servers or other network devices may be stored locally and/or remotely, as appropriate. Where a system includes computerized devices, each such device can
20 include hardware elements that may be electrically coupled via a bus, the elements including, for example, at least one central processing unit (CPU), at least one input device (*e.g.*, a mouse, keyboard, controller, touch-sensitive display element or keypad) and at least one output device (*e.g.*, a display device, printer or speaker). Such a system may also include one or more storage devices, such as disk drives, optical storage
25 devices and solid-state storage devices such as random access memory (RAM) or read-only memory (ROM), as well as removable media devices, memory cards, flash cards, etc.

[0050] Such devices can also include a computer-readable storage media reader, a communications device (*e.g.*, a modem, a network card (wireless or wired), an infrared
30 communication device) and working memory as described above. The computer-readable storage media reader can be connected with, or configured to receive, a computer-readable storage medium representing remote, local, fixed and/or removable

storage devices as well as storage media for temporarily and/or more permanently containing, storing, transmitting and retrieving computer-readable information. The system and various devices also typically will include a number of software applications, modules, services or other elements located within at least one working
5 memory device, including an operating system and application programs such as a client application or Web browser. It should be appreciated that alternate embodiments may have numerous variations from that described above. For example, customized hardware might also be used and/or particular elements might be implemented in hardware, software (including portable software, such as applets) or both. Further,
10 connection to other computing devices such as network input/output devices may be employed.

[0051] Storage media and computer readable media for containing code, or portions of code, can include any appropriate media known or used in the art, including storage media and communication media, such as but not limited to volatile and non-volatile,
15 removable and non-removable media implemented in any method or technology for storage and/or transmission of information such as computer readable instructions, data structures, program modules or other data, including RAM, ROM, EEPROM, flash memory or other memory technology, CD-ROM, digital versatile disk (DVD) or other optical storage, magnetic cassettes, magnetic tape, magnetic disk storage or other
20 magnetic storage devices or any other medium which can be used to store the desired information and which can be accessed by a system device. Based on the disclosure and teachings provided herein, a person of ordinary skill in the art will appreciate other ways and/or methods to implement the various embodiments.

[0052] The specification and drawings are, accordingly, to be regarded in an
25 illustrative rather than a restrictive sense. It will, however, be evident that various modifications and changes may be made thereunto without departing from the broader spirit and scope of the invention as set forth in the claims.

Clauses:

1. A method of rectifying a set of cameras, comprising:
30 receiving an instruction to rectify a set of cameras of a computing device, the set of cameras arranged in a rectangular array wherein pairs of the set of cameras are operable

- to capture three-dimensional images, the set of cameras including at least four cameras positioned on corners of a rectangle of the rectangular array;
- capturing a respective first image using each camera of the set of cameras;
- identifying a plurality of feature points corresponding to one or more objects
- 5 represented in each respective first image;
- determining coordinates for the feature points in each respective first image;
- determining calibration parameters to correct for yaw, pitch, and roll misalignment for each camera of the set of cameras by processing the coordinates for the feature points in each first image;
- 10 storing the calibration parameters for the set of cameras;
- capturing a respective second image using at least two cameras of the set of cameras;
- and
- using the calibration parameters for the at least two cameras to construct a three-dimensional image for the respective second image.
- 15 2. The method of clause 1, wherein the instruction is received at most a determined amount of time after a previous instruction, wherein a small angle approximation is available for use in determining the calibration parameters.
3. The method of clause 2, further comprising:
- generating a set of linear equations for the set of cameras, the linear equations including
- 20 variables for misalignment in yaw, pitch, and roll for each camera of the set of cameras;
- and
- solving the set of linear equations to obtain a value for an amount of yaw, pitch, and roll misalignment for each camera of the set of cameras.
4. The method of clause 1, further comprising:
- 25 initializing a homography for each camera of the set of cameras;
- for each selected camera of the set of cameras, holding homographies for remaining cameras in the set of cameras fixed and matching the homography for the selected camera to the homographies for the remaining cameras; and

using data from the homographies to correct for the yaw, pitch, and roll misalignment for each camera of the set of cameras.

5. The method of clause 1, wherein the calibration parameters are further determined using coordinates for the feature points in at least one additional respective
5 image captured by each camera of the set of cameras.

6. A computer-implemented method, comprising:

obtaining image data captured by at least four cameras, each camera of the at least four cameras capturing a respective image corresponding to a portion of the image data, at least two pairs of the at least four cameras being arranged along orthogonal axes;

- 10 determining coordinates for each of a plurality of feature points as represented in each of the respective images;

using the coordinates associated with each respective image to determine a misalignment of at least one camera of the at least four cameras in at least one of three dimensions, the misalignment being determined by considering respective images

- 15 captured by pairs of the at least four cameras; and

storing information for the misalignment for use in correcting for misalignment errors for at least one subsequent image captured by at least one camera of the at least four cameras.

7. The computer-implemented method of clause 6, further comprising:

- 20 for respective images analyzed for each pair of cameras along one of the orthogonal axes, determining an offset of the feature points along an axis of the orthogonal axes.

8. The computer-implemented method of clause 6, further comprising:

determining a set of linear equations for the at least four cameras, wherein determining the misalignment of at least one camera of the at least four cameras in at least one of

- 25 three dimensions includes solving the set of linear equations, the linear equations including variables for coordinates of the feature points in the respective images captured by each camera of the at least four cameras.

9. The computer-implemented method of clause 6, wherein the misalignment is determined using small angle approximations.

10. The computer-implemented method of clause 6, further comprising:
initializing a homography for each camera of the at least four cameras;
for each selected camera of the at least four cameras, holding homographies for
remaining cameras in the at least four cameras fixed and matching the homography for
5 the selected camera to the homographies for the remaining cameras; and
using the homographies to correct for the misalignment.
11. The computer-implemented method of clause 10, further comprising:
repeating the holding and matching steps for the homography for each of the at least
four cameras until a cost function is determined to converge.
- 10 12. The computer-implemented method of clause 6, wherein the respective image
data captured by the at least four cameras are acquired at substantially the same time.
13. The computer-implemented method of clause 6, wherein the three dimensions
include roll, pitch, and yaw.
14. The computer-implemented method of clause 6, further comprising:
15 determining each of the plurality of feature points using a feature detection algorithm;
and
locating corresponding points of the plurality of feature points in each of the respective
images.
15. The computer-implemented method of clause 6, further comprising:
20 using hypothesis testing to select values for the misalignment, of at least one of the at
least four cameras in at least one of three dimensions, as determined using one or more
non-derivative based cost minimization algorithms.
16. A computing device, comprising:
at least one processor;
25 four cameras arranged along two orthogonal axes, the four cameras having substantially
parallel optical axes; and
memory including instructions that, when executed by the at least one processor, enable
the computing device to:

acquire one or more corresponding images using each camera of the four cameras;
determine coordinates for each of a plurality of feature points as represented in each of the corresponding images;
determine a misalignment of at least one camera of the cameras in at least one of three
5 directions based at least in part upon the coordinates associated with each of the corresponding images, the misalignment being determined by considering corresponding images for pairs of the cameras oriented along at least one of the orthogonal axes; and
store information for the misalignment for use in correcting for misalignment errors in
10 at least one subsequent image captured by at least one of the cameras.

17. The computing device of clause 16, wherein the instructions when executed further cause the computing device to:

determine a set of linear equations for the cameras, wherein determining the misalignment of at least one of the cameras in at least one of three dimensions includes
15 solving the set of linear equations, the linear equations including variables for coordinates of the feature points in the corresponding images.

18. The computing device of clause 16, wherein the instructions when executed further cause the computing device to:

initialize a homography for each camera of the cameras;
20 for each selected camera of the four cameras, hold homographies for remaining cameras of the four cameras fixed and match the homography for the selected camera to the homographies for the remaining cameras; and
accept data from the homographies for use in correcting for the misalignment.

19. The computing device of clause 16, wherein the instructions when executed
25 further cause the computing device to:

use hypothesis testing to select values for the misalignment, of at least one camera of the at least four cameras in at least one of three dimensions, as determined using one or more non-derivative based cost minimization algorithms.

20. The computing device of clause 16, further comprising:

a motion sensor, the computing device configured to wait to automatically acquire the corresponding images until the computing device is substantially motionless as determined using the motion sensor.

WHAT IS CLAIMED IS:

- 1 1. A method of rectifying a set of cameras, comprising:
2 receiving an instruction to rectify a set of cameras of a computing device,
3 the set of cameras arranged in a rectangular array wherein pairs of the set of cameras are
4 operable to capture three-dimensional images, the set of cameras including at least four
5 cameras positioned on corners of a rectangle of the rectangular array;
6 capturing a respective first image using each camera of the set of
7 cameras;
8 identifying a plurality of feature points corresponding to one or more
9 objects represented in each respective first image;
10 determining coordinates for the feature points in each respective first
11 image;
12 determining calibration parameters to correct for yaw, pitch, and roll
13 misalignment for each camera of the set of cameras by processing the coordinates for
14 the feature points in each first image;
15 storing the calibration parameters for the set of cameras;
16 capturing a respective second image using at least two cameras of the set
17 of cameras; and
18 using the calibration parameters for the at least two cameras to construct
19 a three-dimensional image for the respective second image.
- 1 2. The method of claim 1, wherein the instruction is received at
2 most a determined amount of time after a previous instruction, wherein a small angle
3 approximation is available for use in determining the calibration parameters.
- 1 3. The method of claim 2, further comprising:
2 generating a set of linear equations for the set of cameras, the linear
3 equations including variables for misalignment in yaw, pitch, and roll for each camera
4 of the set of cameras; and
5 solving the set of linear equations to obtain a value for an amount of
6 yaw, pitch, and roll misalignment for each camera of the set of cameras.

1 4. The method of claim 1, further comprising:
2 initializing a homography for each camera of the set of cameras;
3 for each selected camera of the set of cameras, holding homographies for
4 remaining cameras in the set of cameras fixed and matching the homography for the
5 selected camera to the homographies for the remaining cameras; and
6 using data from the homographies to correct for the yaw, pitch, and roll
7 misalignment for each camera of the set of cameras.

1 5. The method of claim 1, wherein the calibration parameters are
2 further determined using coordinates for the feature points in at least one additional
3 respective image captured by each camera of the set of cameras.

1 6. A computer-implemented method, comprising:
2 obtaining image data captured by at least four cameras, each camera of
3 the at least four cameras capturing a respective image corresponding to a portion of the
4 image data, at least two pairs of the at least four cameras being arranged along
5 orthogonal axes;
6 determining coordinates for each of a plurality of feature points as
7 represented in each of the respective images;
8 using the coordinates associated with each respective image to determine
9 a misalignment of at least one camera of the at least four cameras in at least one of three
10 dimensions, the misalignment being determined by considering respective images
11 captured by pairs of the at least four cameras; and
12 storing information for the misalignment for use in correcting for
13 misalignment errors for at least one subsequent image captured by at least one camera
14 of the at least four cameras.

1 7. The computer-implemented method of claim 6, further
2 comprising:
3 for respective images analyzed for each pair of cameras along one of the
4 orthogonal axes, determining an offset of the feature points along an axis of the
5 orthogonal axes.

1 8. The computer-implemented method of claim 6, further
2 comprising:
3 determining a set of linear equations for the at least four cameras,
4 wherein determining the misalignment of at least one camera of the at least four
5 cameras in at least one of three dimensions includes solving the set of linear equations,
6 the linear equations including variables for coordinates of the feature points in the
7 respective images captured by each camera of the at least four cameras.

1 9. The computer-implemented method of claim 6, further
2 comprising:
3 initializing a homography for each camera of the at least four cameras;
4 for each selected camera of the at least four cameras, holding
5 homographies for remaining cameras in the at least four cameras fixed and matching the
6 homography for the selected camera to the homographies for the remaining cameras;
7 and
8 using the homographies to correct for the misalignment.

1 10. The computer-implemented method of claim 9, further
2 comprising:
3 repeating the holding and matching steps for the homography for each of
4 the at least four cameras until a cost function is determined to converge.

1 11. A computing device, comprising:
2 at least one processor;
3 four cameras arranged along two orthogonal axes, the four cameras
4 having substantially parallel optical axes; and
5 memory including instructions that, when executed by the at least one
6 processor, enable the computing device to:
7 acquire one or more corresponding images using each camera of
8 the four cameras;
9 determine coordinates for each of a plurality of feature points as
10 represented in each of the corresponding images;
11 determine a misalignment of at least one camera of the cameras in
12 at least one of three directions based at least in part upon the coordinates

13 associated with each of the corresponding images, the misalignment being
14 determined by considering corresponding images for pairs of the cameras
15 oriented along at least one of the orthogonal axes; and
16 store information for the misalignment for use in correcting for
17 misalignment errors in at least one subsequent image captured by at least one of
18 the cameras.

1 12. The computing device of claim 11, wherein the instructions when
2 executed further cause the computing device to:
3 determine a set of linear equations for the cameras, wherein determining
4 the misalignment of at least one of the cameras in at least one of three dimensions
5 includes solving the set of linear equations, the linear equations including variables for
6 coordinates of the feature points in the corresponding images.

1 13. The computing device of claim 11, wherein the instructions when
2 executed further cause the computing device to:
3 initialize a homography for each camera of the cameras;
4 for each selected camera of the four cameras, hold homographies for
5 remaining cameras of the four cameras fixed and match the homography for the selected
6 camera to the homographies for the remaining cameras; and
7 accept data from the homographies for use in correcting for the
8 misalignment.

1 14. The computing device of claim 11, wherein the instructions when
2 executed further cause the computing device to:
3 use hypothesis testing to select values for the misalignment, of at least
4 one camera of the at least four cameras in at least one of three dimensions, as
5 determined using one or more non-derivative based cost minimization algorithms.

1 15. The computing device of claim 11, further comprising:
2 a motion sensor, the computing device configured to wait to
3 automatically acquire the corresponding images until the computing device is
4 substantially motionless as determined using the motion sensor.




FIG. 1




FIG. 2




FIG. 3(a)




FIG. 3(b)




FIG. 3(c)




FIG. 3(d)




FIG. 4(a)




FIG. 4(b)




FIG. 4(c)




FIG. 4(d)




FIG. 4(e)




FIG. 5




FIG. 6




FIG. 7




FIG. 8

INTERNATIONAL SEARCH REPORT

International application No.

PCT/US2014/033401

A. CLASSIFICATION OF SUBJECT MATTER

IPC(8) - G06K 9/46 (2014.01)

USPC - 382/154

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)

IPC(8) - G06K 9/46; G06T 7/20, 15/00; H04N 13/02 (2014.01)

USPC - 348/42, 348/47; 382/154, 382/294

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched
CPC - G06K 9/46; G06T 7/20, 15/00 (2014.02)

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

PatBase, Google Patents, ProQuest

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	WO 2012/044216 A1 (MODEN) 05 April 2012 (05.04.2012) entire document.	1-15
Y	US 2007/0126334 A1 (NAKAMURA et al) 07 June 2007 (07.06.2007) entire document.	1-15
Y	US 2010/0166294 A1 (MARRION et al) 01 July 2010 (01.07.2010) entire document.	1-5
Y	US 5,606,627 A (KUO) 25 February 1997 (25.02.1997) entire document.	2, 3
Y	US 2003/0156751 A1 (LEE et al) 21 August 2003 (21.08.2003) entire document.	4, 9, 10, 13
Y	US 2012/0275667 A1 (LU) 01 November 2012 (01.11.2012) entire document.	6-15
A	US 2012/0237114 A1 (PARK et al) 20 September 2012 (20.09.2012) entire document.	1-15
A	US 2013/0002814 A1 (PARK et al) 03 January 2013 (03.01.2013) entire document.	1-15
A	US 2011/0128353 A1 (LEUNG et al) 02 June 2011 (02.06.2011) entire document.	1-15
A	US 2008/0002879 A1 (JEON et al) 03 January 2008 (03.01.2008) entire document.	1-15
A	US 2007/0165942 A1 (JIN et al) 19 July 2007 (19.07.2007) entire document.	1-15

☐ Further documents are listed in the continuation of Box C.

* 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

25 July 2014

Date of mailing of the international search report

18 AUG 2014

Name and mailing address of the ISA/US

Mail Stop PCT, Attn: ISA/US, Commissioner for Patents
P.O. Box 1450, Alexandria, Virginia 22313-1450

Facsimile No. 571-273-3201

Authorized officer:

Blaine R. Copenheaver

PCT Helpdesk: 571-272-4300

PCT OSP: 571-272-7774