




(43) International Publication Date
21 September 2017 (21.09.2017)

- (51) International Patent Classification:
H04N 13/02 (2006.01) *G06T 7/80* (2017.01)
- (21) International Application Number:
PCT/US2017/014825
- (22) International Filing Date:
25 January 2017 (25.01.2017)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:
15/074,166 18 March 2016 (18.03.2016) US
- (71) Applicant: INTEL CORPORATION [US/US]; 2200 Mission College Blvd., Santa Clara, California 95054 (US).
- (72) Inventor: ELDAR, Avigdor; 15 Steinberg Street, 96771 Jerusalem, JM (IL).
- (74) Agents: HOLOMON, Jamilla N. et al.; INTERNATIONAL IP LAW GROUP, P.L.L.C., 13231 Champion Forest Dr., Suite 410, Houston, Texas 77069 (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, DJ, 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, KH, KN, KP, KR, KW, 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).

[Continued on next page]

(54) Title: CAMERA CALIBRATION USING DEPTH DATA



(57) Abstract: An apparatus is described herein. The apparatus includes an image capture module to capture depth data and sensor data for a plurality of views and an extraction module to extract a first plurality of features from the depth data and a second plurality of features from the sensor data for each view. The apparatus also includes a correspondence module to locate corresponding features in the first plurality of features and the second plurality of features for each view and a depth module to generate three-dimensional data for each feature of the first plurality of features for each view. Additionally, the apparatus includes a calibration module to calibrate the multiple cameras by matching the generated three dimensional data with the corresponding features in the first plurality of features and the second plurality of features.

400
FIG. 4A

**Declarations under Rule 4.17:**

- *as to the identity of the inventor (Rule 4.17(i))*
- *as to applicant's entitlement to apply for and be granted a patent (Rule 4.17(ii))*

Published:

- *with international search report (Art. 21(3))*

CAMERA CALIBRATION USING DEPTH DATA

CROSS REFERENCE TO RELATED APPLICATION

[0001] The present application claims the benefit of the filing date of United States Patent Application Serial No. 15/074,166, by Avigdor Eldar, entitled "Camera Calibration Using Depth Data," filed March 18, 2016, and is incorporated herein by reference.

BACKGROUND ART

[0002] Electronic devices, such as tablets, phablets, smartphones, mobile phones, desktops, laptops, gaming devices, all-in-one systems, and the like may include various cameras for capturing images. The electronic devices may require calibration of camera with regard to how the camera perceives the world. In some cases, the electronic device may include a plurality of cameras configured to work individually or in concert. Each camera may be used to capture different image information, such as depth, color space, or other information regarding a scene.

BRIEF DESCRIPTION OF THE DRAWINGS

[0003] Fig. 1 is a block diagram of an exemplary system that enables camera calibration using depth data;

[0004] Fig. 2A is a diagram illustrating a target to be used in calibration;

[0005] Fig. 2B is an illustration of a target placed in error for use in calibration;

[0006] Fig. 3 is an illustration of a plurality of views of a calibration target;

[0007] Fig. 4A is a process flow diagram of a method that enables camera calibration using depth data;

[0008] Fig. 4B is a process flow diagram of another method that enables camera calibration using depth data;

[0009] Fig. 5 is a line graph illustrating error with traditional calibration and averaged calibration; and

[0010] Fig. 6 is a block diagram showing media that contains logic for camera calibration using depth data.

[0011] The same numbers are used throughout the disclosure and the figures to reference like components and features. Numbers in the 100 series refer to features

originally found in Fig. 1; numbers in the 200 series refer to features originally found in Fig. 2; and so on.

DETAILED DESCRIPTION

[0012] Typically, each camera of a plurality of cameras in a system is calibrated prior to use. In some cases, the calibration of multiple cameras require capturing many images of a calibration target at different angles and distances from the cameras of a given system that includes multiple cameras. The calibration target may be a pattern with an easily distinguished geometry. For example, the calibration target may be a printed checkerboard. In the typical calibration process, the checkerboard must have a precise and accurate physical geometry. There is often high costs associated with generating a precise and accurate calibration target.

[0013] Embodiments described herein enable camera calibration using depth data. In embodiments, the present techniques enable a depth camera to red, green, blue (RGB) camera calibration. Through this calibration, the depth camera and RGB camera may be set such that the cameras properly capture images. Calibration may be performed from a depth sensor to an RGB sensor, utilizing available three dimensional depth data instead of limited two dimensional image data. The actual physical target geometry is computed as part of the process, rather than required as input to the calibration process.

[0014] Some embodiments may be implemented in one or a combination of hardware, firmware, and software. Further, some embodiments may also be implemented as instructions stored on a machine-readable medium, which may be read and executed by a computing platform to perform the operations described herein. A machine-readable medium may include any mechanism for storing or transmitting information in a form readable by a machine, e.g., a computer. For example, a machine-readable medium may include read only memory (ROM); random access memory (RAM); magnetic disk storage media; optical storage media; flash memory devices; or electrical, optical, acoustical or other form of propagated signals, e.g., carrier waves, infrared signals, digital signals, or the interfaces that transmit and/or receive signals, among others.

[0015] An embodiment is an implementation or example. Reference in the specification to “an embodiment,” “one embodiment,” “some embodiments,” “various embodiments,” or “other embodiments” means that a particular feature, structure, or

characteristic described in connection with the embodiments is included in at least some embodiments, but not necessarily all embodiments, of the present techniques. The various appearances of “an embodiment,” “one embodiment,” or “some embodiments” are not necessarily all referring to the same embodiments. Elements or aspects from an embodiment can be combined with elements or aspects of another embodiment.

[0016] Not all components, features, structures, characteristics, etc. described and illustrated herein need be included in a particular embodiment or embodiments. If the specification states a component, feature, structure, or characteristic “may”, “might”, “can” or “could” be included, for example, that particular component, feature, structure, or characteristic is not required to be included. If the specification or claim refers to “a” or “an” element, that does not mean there is only one of the element. If the specification or claims refer to “an additional” element, that does not preclude there being more than one of the additional element.

[0017] It is to be noted that, although some embodiments have been described in reference to particular implementations, other implementations are possible according to some embodiments. Additionally, the arrangement and/or order of circuit elements or other features illustrated in the drawings and/or described herein need not be arranged in the particular way illustrated and described. Many other arrangements are possible according to some embodiments.

[0018] In each system shown in a figure, the elements in some cases may each have a same reference number or a different reference number to suggest that the elements represented could be different and/or similar. However, an element may be flexible enough to have different implementations and work with some or all of the systems shown or described herein. The various elements shown in the figures may be the same or different. Which one is referred to as a first element and which is called a second element is arbitrary.

[0019] Fig. 1 is a block diagram of an exemplary system that enables camera calibration using depth data. The electronic device 100 may be, for example, a laptop computer, tablet computer, mobile phone, smart phone, or a wearable device, among others. The electronic device 100 may include a central processing unit (CPU) 102 that is configured to execute stored instructions, as well as a memory device 104 that stores instructions that are executable by the CPU 102. The CPU may be coupled to the memory device 104 by a bus 106. Additionally, the CPU 102

can be a single core processor, a multi-core processor, a computing cluster, or any number of other configurations. Furthermore, the electronic device 100 may include more than one CPU 102. The memory device 104 can include random access memory (RAM), read only memory (ROM), flash memory, or any other suitable memory systems. For example, the memory device 104 may include dynamic random access memory (DRAM).

[0020] The electronic device 100 also includes a graphics processing unit (GPU) 108. As shown, the CPU 102 can be coupled through the bus 106 to the GPU 108. The GPU 108 can be configured to perform any number of graphics operations within the electronic device 100. For example, the GPU 108 can be configured to render or manipulate graphics images, graphics frames, videos, streaming data, or the like, to be rendered or displayed to a user of the electronic device 100. In some embodiments, the GPU 108 includes a number of graphics engines, wherein each graphics engine is configured to perform specific graphics tasks, or to execute specific types of workloads.

[0021] The CPU 102 can be linked through the bus 106 to a display interface 110 configured to connect the electronic device 100 to one or more display devices 112. The display devices 112 can include a display screen that is a built-in component of the electronic device 100. In embodiments, the display interface 110 is coupled with the display devices 112 via any networking technology such as cellular hardware 124, Wifi hardware 126, or Bluetooth Interface 128 across the network 132. The display devices 112 can also include a computer monitor, television, or projector, among others, that is externally connected to the electronic device 100.

[0022] The CPU 102 can also be connected through the bus 106 to an input/output (I/O) device interface 114 configured to connect the electronic device 100 to one or more I/O devices 116. The I/O devices 116 can include, for example, a keyboard and a pointing device, wherein the pointing device can include a touchpad or a touchscreen, among others. The I/O devices 116 can be built-in components of the electronic device 100, or can be devices that are externally connected to the electronic device 100. Accordingly, in embodiments, the I/O device interface 114 is coupled with the I/O devices 116 via any networking technology such as cellular hardware 124, Wifi hardware 126, or Bluetooth Interface 128 across the network 132. The I/O devices 116 can also include any I/O device that is externally connected to the electronic device 100.

[0023] The electronic device 100 also includes image capture mechanisms 118. The image capture mechanisms may be a plurality of cameras. The image capture mechanisms 118 may also include a plurality of sensors. In embodiments, the image capture mechanisms 118 may be a depth camera and an RGB camera. Additionally, in embodiments, the image capture mechanisms 118 may be a depth sensor and an RGB sensor. In some embodiments, the image capture mechanisms 118 can be a camera, stereoscopic camera, infrared sensor, and the like. The image capture mechanisms 118 are used to capture image information and the corresponding depth information. The image capture mechanisms 118 may include sensors such as a depth sensor, RGB sensor, an image sensor, an infrared sensor, an X-Ray photon counting sensor, a light sensor, or any combination thereof. The image sensors may include charge-coupled device (CCD) image sensors, complementary metal-oxide-semiconductor (CMOS) image sensors, system on chip (SOC) image sensors, image sensors with photosensitive thin film transistors, or any combination thereof. In some embodiments, a sensor is a depth sensor. The depth sensor may be used to capture the depth information associated with the image information. In some embodiments, a driver may be used to operate a sensor within the image capture devices 118, such as a depth sensor.

[0024] The electronic device 118 may also include a calibration mechanism 120. The calibration mechanism may use depth data to make the calibration process more robust. In embodiments, the calibration mechanism 120 is to calibrate a plurality of cameras using depth data, thereby removing the burden of having accurate targets and potentially introduce new calibration scenarios. New calibration scenarios include, but are not limited to natural objects as targets instead of designated checkerboard targets. During calibration, the depth camera center is maintained at a center location relative to the calibration target. For example, the location may be described according to Cartesian coordinates, where x , y , and z are each zero. This single frame of reference is used to capture all views of the calibration target onto a reference location to create an average calibration target model. In embodiments, a calibration target model is a function used to transformation between world and image coordinates. Calibration is used to obtain the correct model parameters.

[0025] In embodiments, camera calibration results in a linear calibration target model that is defined by eleven parameters. The eleven parameters include camera

location (3 parameters: x, y, z), orientation (3 parameters: roll, pitch, and yaw), focal length (1 parameter: distance), pixel scale (1 parameter), pixel aspect ratio (1 parameter), and image plane center offset (2 parameters). The calibration target model, through these parameters projection for each camera that maps a point in the camera image from a three-dimensional point in space to a two-dimensional location in the camera image.

[0026] The electronic device 100 also includes a storage device 124. The storage device 124 is a physical memory such as a hard drive, an optical drive, a flash drive, an array of drives, or any combinations thereof. The storage device 124 can store user data, such as audio files, video files, audio/video files, and picture files, among others. The storage device 124 can also store programming code such as device drivers, software applications, operating systems, and the like. The programming code stored to the storage device 124 may be executed by the CPU 102, GPU 108, or any other processors that may be included in the electronic device 100.

[0027] The CPU 102 may be linked through the bus 106 to cellular hardware 126. The cellular hardware 126 may be any cellular technology, for example, the 4G standard (International Mobile Telecommunications-Advanced (IMT-Advanced) Standard promulgated by the International Telecommunications Union – Radio communication Sector (ITU-R)). In this manner, the electronic device 100 may access any network 132 without being tethered or paired to another device, where the cellular hardware 126 enables access to the network 132.

[0028] The CPU 102 may also be linked through the bus 106 to WiFi hardware 128. The WiFi hardware 128 is hardware according to WiFi standards (standards promulgated as Institute of Electrical and Electronics Engineers' (IEEE) 802.11 standards). The WiFi hardware 128 enables the electronic device 100 to connect to the Internet using the Transmission Control Protocol and the Internet Protocol (TCP/IP). Accordingly, the electronic device 100 can enable end-to-end connectivity with the Internet by addressing, routing, transmitting, and receiving data according to the TCP/IP protocol without the use of another device. Additionally, a Bluetooth Interface 130 may be coupled to the CPU 102 through the bus 106. The Bluetooth Interface 130 is an interface according to Bluetooth networks (based on the Bluetooth standard promulgated by the Bluetooth Special Interest Group). The Bluetooth Interface 130 enables the electronic device 100 to be paired with other

Bluetooth enabled devices through a personal area network (PAN). Accordingly, the network 132 may be a PAN. Examples of Bluetooth enabled devices include a laptop computer, desktop computer, ultrabook, tablet computer, mobile device, or server, among others.

[0029] The block diagram of Fig. 1 is not intended to indicate that the electronic device 100 is to include all of the components shown in Fig. 1. Rather, the computing system 100 can include fewer or additional components not illustrated in Fig. 1 (e.g., sensors, power management integrated circuits, additional network interfaces, etc.). The electronic device 100 may include any number of additional components not shown in Fig. 1, depending on the details of the specific implementation. Furthermore, any of the functionalities of the CPU 102 may be partially, or entirely, implemented in hardware and/or in a processor. For example, the functionality may be implemented with an application specific integrated circuit, in logic implemented in a processor, in logic implemented in a specialized graphics processing unit, or in any other device.

[0030] Fig. 2A is a diagram illustrating a target 200A to be used in calibration. As illustrated, the calibration target 200A is a checkerboard pattern. Thus, the target 200A includes a plurality of squares, where the plurality of squares include a plurality of black squares 202 and a plurality of white squares 204. In examples, for use as a calibration target, each square is 2 centimeters by 2 centimeters.

[0031] In existing calibration processes, the physical geometry of the target such as the target 200A must be known in advance and is an input to a calibration algorithm. The physical geometry may include either three dimensional coordinates for a three dimensional target or two dimensional measurements for the more commonly used, two dimensional, planar targets. In both, cases achieving an accurate target model is an expensive process. For example, printing a checkerboard pattern for a two dimensional calibration target with squares that are precisely a particular size can be an expensive process, and requires costly equipment to generate a precise and accurate target.

[0032] Additionally, an accurate and precisely printed target may be incorrectly placed for calibration by an operator. Fig. 2B is an illustration of a target 200B placed in error for use in calibration. For calibration, an operator typically places the calibration target on a platform so that cameras of the electronic device can be calibrated by capturing the calibration target using the cameras. The calibration

target may be printed onto paper that is applied to a platform using an adhesive. The target 200B illustrates deformities that can occur when placing the calibration target. For example, bends 206 and 208 illustrate deformities that may be present when the calibration target is placed. In some cases, the bends 206 and 208 are a result of dimpling, puckering, or an otherwise non-smooth placement of the checkerboard target onto a platform. The bends 206 and 208 result in non-planar regions that are not expected in a typical calibration scenario. If these deformities are present when using the conventional calibration that requires prior knowledge of the size and location of the calibration target, conventional calibration may fail or provide erroneous results.

[0033] Thus, to implement conventional calibration the calibration target must be accurately and precisely printed with extremely tight tolerances to satisfy a calibration algorithm that requires these exact measurements. The present techniques overcome the need for a precise and accurate calibration target by using depth data to reconstruct the calibration target. In this manner, the exact dimensions of the calibration target as discovered, thereby eliminating the need for an accurately and precisely printed calibration target.

[0034] Since the calibration target does not need to be known in advance, the calibration according to the present techniques can be implemented in the field because the calibration target does not need to be known in advance. As discussed above, the calibration process described herein utilizes available depth data to calculate an accurate and precise three dimensional geometry of the actual physical structure of the calibration target. In the manner, accurate and precise target creation is eliminated. Moreover, the present techniques also reduces the setup burden of the operator.

[0035] Depth data, as used herein, may refer to representations of depth information such as a depth field, a point cloud, a depth map, or a three dimensional (three dimensional) polygonal mesh that may be used to indicate the depth of three dimensional objects within the image. While the techniques are described herein using a depth data, any depth representation can be used. In some cases, depth data used to reconstruct the calibration target can be noisy. To overcome noisy depth data, the present techniques employ smoothing techniques. In embodiments, smoothing may be performed via averaging as described herein as well as sub-pixel refinement.

[0036] Fig. 3 is an illustration of a plurality of views 300 of a calibration target. The camera center is located at reference number 302. The plurality of views includes a view 304, a view 308, and a view 310. For ease of description, three views are illustrated. However, the present techniques can be implemented with any number of views.

[0037] The camera may be placed in a single frame of reference with regard to each view. In embodiments, the single frame of reference, as described herein, indicates a frame reference relative to the camera. Specifically, the depth camera center may be located at $x, y, z = (0,0,0)$ for all three dimensional views 304, 306, and 308 of the target. Each view may be captured by repositioning the calibration target or the camera. Once a number of views are captured, the multiple views may be aligned onto a reference location.

[0038] Put another way, as a result of the depth data obtained from the depth camera, for each capture of the calibration target, the three dimensional location (X,Y,Z) relative to the camera can be derived. The camera may be positioned at coordinate $x, y, z = (0,0,0)$ in space, and the three dimensional coordinates of all views can be derived using the depth data from each view. In other words, the camera acts as the single frame of reference for all views. In embodiments, the present techniques may use multiple reference locations. For example, the corners of each view may be used to define reference locations for each view.

[0039] After the multiple views have been aligned onto a reference location, the views may be averaged to create an average, less noisy, "model." Thus, in each of views 304, 306, and 308, there is a common point that provides the same view in each of the views. Features of the calibration target may be found in each view, and then a transform is found for each view. A close formula may be used to estimate a rigid transform (camera pose) between two locations of the same object in space. This technique may be used when noise (or measurement error) is minimal. In embodiments, a Best Rigid Transform is used to obtain a transform for each view. Alternatively, the depth data from the multiple views can be averaged by tracking the camera pose as each view is captured. The camera pose, and how each camera pose relates to the other tracked camera poses can be used to average the views into a single model. In embodiments, when the depth data is noisy, a Kinect Fusion technique may be used to detect a rigid transform for each view.

[0040] With conventional calibration techniques, averaging is not possible because there is no unique reference frame. In essence, there is no equivalent averaging process for two dimensional measurements as this is the calibration process itself. Put another way, when using a single set of measurements in a single dimension, any attempt to find the relation between two dimensional images of the target results in a three dimensional to two dimensional projection transformation, which requires knowledge of camera intrinsic parameters.

[0041] Fig. 4A is a process flow diagram of a method 400 that enables camera calibration using depth data. At block 402, stereoscopic image pairs are captured. The stereoscopic image pairs may be captured for a plurality of views. The depth data may be captured as a stereoscopic infrared (IR) image pair. An RGB image may also be captured for each view. In embodiments, an image set includes an RGB image and a depth image. In examples, features may be detected from an image set including RGB and IR images and matched. In this example, the depth and resulting 3D coordinate may be computed for each matched pair. In another example, the image set can include structured light/coded light, where the input is a depth map and a single RGB/IR image. Features may be detected on IR/ RGB image, the corresponding 3D coordinate is then computed from depth map.

[0042] At block 404, features are extracted from each image. Correspondence for each image set may be found for each feature in the RGB image and the depth image. As used herein, features may be points of interest on a calibration target. For example, with a checkerboard, the intersection of each square with another square may be a feature. Features may also include edges, lines, corners, and the like.

[0043] Thus, for each real world coordinate, there are two image coordinates (one derived from the depth image pair and one in the RGB image) that correspond to the real world coordinate. At block 406, three dimensional data is generated using the image pairs. In embodiments, a depth module is used to generate three dimensional data from the captured depth image. Stereo image pairs may be used to explicitly reconstruct per feature three dimensional data.

[0044] In embodiments, subpixel refinement is performed for each detected feature. Subpixel refinement may act as a function to compute a more accurate image location of the detected feature. By observing that each vector from the center the current pixel to a second point located within a neighborhood of the pixel

is orthogonal to the image gradient at subject to image and measurement noise. A sum of the gradients in the neighborhood of the pixel results in a new, more accurate pixel location. This process can be performed iteratively until the coordinate location stays within a particular threshold.

[0045] At block 408, the three dimensional data is refined by averaging all views. Averaging may be used to reduce the noise that may occur in each individual view, where each view includes a depth image. Because each capture is within a single frame of reference, three dimensional correspondences between the calibration targets can be found and the results can be averaged. In embodiments, the average may be found by finding a best fit rigid transform for each view to a reference view. The transform may be the best rotation and/or translation that will align the points in the current view to the reference view. In embodiments, 'best' is in terms of least square errors. This transform may also be called the Euclidean or Rigid transform since the shape and size of the view is preserved. For each view, a transform is computed such that

$$V' = Tv(V)$$

where Tv is the best fit rigid transform, V is the current view, and V' is very close to the reference view. Each transformed view V' is then averaged, which results in a calibration target model (TM). Now that a calibration target model TM has been found, each view can be replaced with a refined view:

$$V_{refined} = Tv^{-1}(TM)$$

[0046] At block 410, the best three dimensional to two dimensional projection is computed based on the refined three dimensional data and the corresponding two dimensional RGB features. The refined three dimensional data and corresponding two dimensional RGB features are found via the calibration target model. Accordingly, the projection is computed using a function that estimates the intrinsic camera parameters and extrinsic parameters for each of the views.

[0047] In embodiments, the extrinsic parameters may be identical for all views as a result of the single frame of reference. These extrinsic parameters describe the calibration between the RGB and depth sensors. Once the model is developed, the three dimensional data from multiple view results in a linear transformation. Additionally, in embodiments, the calibration results in the capture of eleven parameters that define the relationship between cameras of the device. The eleven

parameters are the intrinsic parameters that define the relation of the RGB camera to the world. Additionally, there are six extrinsic parameters that define the relationship between the RGB and depth sensors (rotation and translation).

[0048] Fig. 4B is a process flow diagram of a method 400B that enables camera calibration using depth data. In Fig. 4B depth data and RGB data are obtained and processed to generate a calibration target model. At block 420, the system is positioned to view the calibration target on all systems. In embodiments, the system may be an electronic device (Fig. 1). At block 422, images are captured using the RGB sensor and the depth sensor. These images may be represented as RGB[i] and Depth[i], where i is the number of views to be obtained. At block 424, features are detected on each of RGB[i] and Depth[i], where the features are represented as RGBfeat[i] and Depthfeat[i]. Each image RGB[i] and Depth[i] may contain a plurality of features.

[0049] At block 426, the features found on RGB[i] are matched to features on Depth[i], and vice versa. The matched features may be represented by Match[i]. At block 428, a location of a three dimensional feature Depthfeat[i], is computed relative to the depth camera. The location of the three dimensional feature may be represented by 3Dfeat[i]. At block 430, the match is applied. In embodiments, the 3Dfeat[i] and Match[i] as used to derive a two dimensional coordinate in the RGB[i] and a corresponding three dimensional coordinate in the Depth[i].

[0050] At block 432, it is determined if i is less than a threshold. If i is less than a threshold, process flow returns to block 420 where the system is positioned to capture another view. If i is greater than a threshold, process flow continues to block 434 sub-pixel refinement is performed. At block 434, each feature in 3Dfeat[i] is processed via sub-pixel refinement. The result of the sub-pixel refinement is an refinement of the three dimensional features across views to obtain a refined pixel location for each feature. At block 436, a refined feature match RefinedFeatMatch is found by using Refined3Dfeat[i] and 3DMatch[i] to derive a two dimensional coordinate in the RGB[i] and a corresponding three dimensional coordinate in the Depth[i].

[0051] At block 438, calibration data is obtained by computing the best three dimensional to two dimensional projection based on the refined three dimensional data Refined3DFeat. In embodiments, an Open Source Computer Vision (OpenCV) calibrate function may be used to obtain the calibration data. OpenCV may refer to

the OpenCV Specification, released December 21, 2015. At block 440, the calibration data may be stored and applied to data captured by the RGB and depth sensors.

[0052] Fig. 5 is a line graph 500 illustrating error with traditional calibration and the calibration techniques described herein. The line graph 500 includes an x-axis 502 representing the number of process images or views. The line graph 500 also includes a y-axis 504 representing errors calibration errors in pixels. The traditional calibration is represented by the line 506, and the averaged calibration is represented by the line 508. As illustrated, the number of calibration errors is lower for the averaged calibration when compared with the traditional calibration, regardless of the number of process views or images.

[0053] As illustrated, the present techniques result in a better absolute calibration error value (approximately 1.6 pixel error vs. approximately 5.4 pixels with the traditional method). In addition, convergence associated with the present techniques is faster, where three images are sufficient to reduce calibration errors, and two images is the absolute minimum required for mathematical correctness.

[0054] The present techniques have better tolerance to operator errors and also employ noise reduction techniques that cannot be employed by traditional calibration methods, neither for 2D nor 3D targets. Moreover, the present techniques have a better resilience to noise, as there are less unknown parameters to compute since the depth camera frame of reference is used for all scenes.

[0055] Fig. 6 is a block diagram showing media 600 that contains logic for camera calibration using depth data. The media 600 may be a computer-readable medium, including a non-transitory medium that stores code that can be accessed by a processor 602 over a computer bus 604. For example, the computer-readable media 600 can be volatile or non-volatile data storage device. The media 600 can also be a logic unit, such as an Application Specific Integrated Circuit (ASIC), a Field Programmable Gate Array (FPGA), or an arrangement of logic gates implemented in one or more integrated circuits, for example.

[0056] The media 600 may include modules 606-614 configured to perform the techniques described herein. For example, an image capture module 606 may be configured to capture images using the depth camera and the RGB camera.

[0057] A correspondence module 608 may be configured to extract features and find the correspondence between features. A depth module 610 may be configured

generate depth data for each feature based on the stereoscopic data captured for each feature point. At block 612, an averaging module averages depth data from all views. At block 614, a calibration module completes calibration by obtaining a three dimensional to two dimensional best fit projection. In some embodiments, the modules 607-614 may be modules of computer code configured to direct the operations of the processor 602.

[0058] The block diagram of Fig. 6 is not intended to indicate that the media 600 is to include all of the components shown in Fig. 6. Further, the media 600 may include any number of additional components not shown in Fig. 6, depending on the details of the specific implementation.

[0059] Example 1 is an apparatus for calibrating multiple cameras. The apparatus includes an image capture module to capture depth data and sensor data for a plurality of views; an extraction module to extract a first plurality of features from the depth data and a second plurality of features from the sensor data for each view; a correspondence module to locate corresponding features in the first plurality of features and the second plurality of features for each view; a depth module to generate three-dimensional data for each feature of the first plurality of features for each view; and a calibration module to calibrate the multiple cameras by matching the generated three dimensional data with the corresponding features in the first plurality of features and the second plurality of features.

[0060] Example 2 includes the apparatus of example 1, including or excluding optional features. In this example, a controller is to average the three-dimensional data across multiple views, and the calibration module is to calibrate the multiple cameras by matching the averaged three dimensional data with the corresponding features in the first plurality of features and the second plurality of features.

[0061] Example 3 includes the apparatus of any one of examples 1 to 2, including or excluding optional features. In this example, the apparatus includes averaging the three dimensional data across multiple views by calculating a best fit rigid transform from each view of a plurality of views to a reference view.

[0062] Example 4 includes the apparatus of any one of examples 1 to 3, including or excluding optional features. In this example, the apparatus includes averaging the three dimensional data by tracking a camera pose as each view of the plurality of views is captured and using the camera pose to average the views into a single view.

[0063] Example 5 includes the apparatus of any one of examples 1 to 4, including or excluding optional features. In this example, the depth data is obtained from a stereoscopic infrared image pair.

[0064] Example 6 includes the apparatus of any one of examples 1 to 5, including or excluding optional features. In this example, the apparatus includes applying subpixel refinement to each feature of the plurality of features.

[0065] Example 7 includes the apparatus of any one of examples 1 to 6, including or excluding optional features. In this example, features are detected by observing points of interest in each view of the plurality of views.

[0066] Example 8 includes the apparatus of any one of examples 1 to 7, including or excluding optional features. In this example, calibrating the multiple views results in a calibration target model used to transformation between world and image coordinates.

[0067] Example 9 includes the apparatus of any one of examples 1 to 8, including or excluding optional features. In this example, the plurality of views include multiple views of a calibration target. Optionally, calibration is performed without prior knowledge of the calibration target.

[0068] Example 10 is a method for calibrating multiple cameras. The method includes capturing depth data and sensor data for a plurality of views; extracting a plurality of corresponding features from the depth data and the sensor data; generating three-dimensional data for each feature of the plurality of corresponding features; and calibrating the multiple cameras by calculating a projection based on the three-dimensional data and the plurality of corresponding features.

[0069] Example 11 includes the method of example 10, including or excluding optional features. In this example, the three-dimensional data is averaged across the multiple views, and the multiple cameras are calibrated by calculating a projection based on the averaged three-dimensional data and the plurality of corresponding features across multiple views.

[0070] Example 12 includes the method of any one of examples 10 to 11, including or excluding optional features. In this example, the method includes averaging the three dimensional data across multiple views by calculating a best fit rigid transform from each view of a plurality of views to a reference view.

[0071] Example 13 includes the method of any one of examples 10 to 12, including or excluding optional features. In this example, the method includes

averaging the three dimensional data by tracking a camera pose as each view of the plurality of views is captured and using the camera pose to average the views into a single view.

[0072] Example 14 includes the method of any one of examples 10 to 13, including or excluding optional features. In this example, the depth data is obtained from a stereoscopic infrared image pair.

[0073] Example 15 includes the method of any one of examples 10 to 14, including or excluding optional features. In this example, the sensor data is an RGB sensor data.

[0074] Example 16 includes the method of any one of examples 10 to 15, including or excluding optional features. In this example, the depth data is obtained from structured light.

[0075] Example 17 includes the method of any one of examples 10 to 16, including or excluding optional features. In this example, the plurality of views include multiple views of a calibration target.

[0076] Example 18 includes the method of any one of examples 10 to 17, including or excluding optional features. In this example, calibration is performed without prior knowledge of a calibration target.

[0077] Example 19 includes the method of any one of examples 10 to 18, including or excluding optional features. In this example, averaging the three dimensional data and applying a subpixel refinement to the plurality of features results in a smoothing of the depth data.

[0078] Example 20 is a system for calibrating multiple cameras. The system includes a depth camera and a sensor to capture images; a memory configured to receive image data; and a processor coupled to the memory, depth camera, and sensor, the processor to: capture depth data and sensor data for a plurality of views; extract a plurality of corresponding features from the depth data and the sensor data; generate three-dimensional data for each feature of the plurality of corresponding features; and calibrate the multiple cameras by calculating a projection based on the three-dimensional data and the plurality of corresponding features.

[0079] Example 21 includes the system of example 20, including or excluding optional features. In this example, the processor is to average the three-dimensional data for each view, and the calibration module is to calibrate the multiple cameras by

matching the averaged three dimensional data with the corresponding features in the first plurality of features and the second plurality of features.

[0080] Example 22 includes the system of any one of examples 20 to 21, including or excluding optional features. In this example, the system includes averaging the three dimensional data across multiple views by calculating a best fit rigid transform from each view of a plurality of views to a reference view.

[0081] Example 23 includes the system of any one of examples 20 to 22, including or excluding optional features. In this example, the system includes averaging the three dimensional data by tracking a camera pose as each view of the plurality of views is captured and using the camera pose to average the views into a single view.

[0082] Example 24 includes the system of any one of examples 20 to 23, including or excluding optional features. In this example, the depth data is obtained from a stereoscopic infrared image pair.

[0083] Example 25 includes the system of any one of examples 20 to 24, including or excluding optional features. In this example, the system includes applying subpixel refinement to each feature of the plurality of features.

[0084] Example 26 includes the system of any one of examples 20 to 25, including or excluding optional features. In this example, features are detected by observing points of interest in each view of the plurality of views.

[0085] Example 27 includes the system of any one of examples 20 to 26, including or excluding optional features. In this example, calibrating the multiple views results in a calibration target model used to transformation between world and image coordinates.

[0086] Example 28 includes the system of any one of examples 20 to 27, including or excluding optional features. In this example, the plurality of views include multiple views of a calibration target. Optionally, calibration is performed without prior knowledge of the calibration target.

[0087] Example 29 is an apparatus for calibrating multiple cameras. The apparatus includes an image capture module to capture depth data and sensor data for a plurality of views; a means to extract a first plurality of features from the depth data and a second plurality of features from the sensor data for each view; a means to locate corresponding features in the first plurality of features and the second plurality of features for each view; a means to generate three-dimensional data for

each feature of the first plurality of features for each view; and a means to calibrate the multiple cameras by matching the generated three dimensional data with the corresponding features in the first plurality of features and the second plurality of features.

[0088] Example 30 includes the apparatus of example 29, including or excluding optional features. In this example, a means to average views is to average the three-dimensional data for each view, and the means to calibrate the multiple cameras is to match the averaged three dimensional data with the corresponding features in the first plurality of features and the second plurality of features.

[0089] Example 31 includes the apparatus of any one of examples 29 to 30, including or excluding optional features. In this example, the apparatus includes averaging the three dimensional data across multiple views by calculating a best fit rigid transform from each view of a plurality of views to a reference view.

[0090] Example 32 includes the apparatus of any one of examples 29 to 31, including or excluding optional features. In this example, the apparatus includes averaging the three dimensional data by tracking a camera pose as each view of the plurality of views is captured and using the camera pose to average the views into a single view.

[0091] Example 33 includes the apparatus of any one of examples 29 to 32, including or excluding optional features. In this example, the depth data is obtained from a stereoscopic infrared image pair.

[0092] Example 34 includes the apparatus of any one of examples 29 to 33, including or excluding optional features. In this example, the apparatus includes applying subpixel refinement to each feature of the plurality of features.

[0093] Example 35 includes the apparatus of any one of examples 29 to 34, including or excluding optional features. In this example, features are detected by observing points of interest in each view of the plurality of views.

[0094] Example 36 includes the apparatus of any one of examples 29 to 35, including or excluding optional features. In this example, calibrating the multiple views results in a calibration target model used to transformation between world and image coordinates.

[0095] Example 37 includes the apparatus of any one of examples 29 to 36, including or excluding optional features. In this example, the plurality of views

include multiple views of a calibration target. Optionally, calibration is performed without prior knowledge of the calibration target.

[0096] Example 38 is at least one machine readable medium comprising a plurality of instructions that. The computer-readable medium includes instructions that direct the processor to capture depth data and sensor data for a plurality of views; extract a plurality of corresponding features from the depth data and the sensor data; generate three-dimensional data for each feature of the plurality of corresponding features; and calibrate the multiple cameras by calculating a projection based on the three-dimensional data and the plurality of corresponding features.

[0097] Example 39 includes the computer-readable medium of example 38, including or excluding optional features. In this example, the three-dimensional data is averaged across the multiple views, and the multiple cameras are calibrated by calculating a projection based on the averaged three-dimensional data and the plurality of corresponding features across multiple views.

[0098] Example 40 includes the computer-readable medium of any one of examples 38 to 39, including or excluding optional features. In this example, the computer-readable medium includes averaging the three dimensional data across multiple views by calculating a best fit rigid transform from each view of a plurality of views to a reference view.

[0099] Example 41 includes the computer-readable medium of any one of examples 38 to 40, including or excluding optional features. In this example, the computer-readable medium includes averaging the three dimensional data by tracking a camera pose as each view of the plurality of views is captured and using the camera pose to average the views into a single view.

[0100] Example 42 includes the computer-readable medium of any one of examples 38 to 41, including or excluding optional features. In this example, the depth data is obtained from a stereoscopic infrared image pair.

[0101] Example 43 includes the computer-readable medium of any one of examples 38 to 42, including or excluding optional features. In this example, the sensor data is an RGB sensor data.

[0102] Example 44 includes the computer-readable medium of any one of examples 38 to 43, including or excluding optional features. In this example, the depth data is obtained from structured light.

[0103] Example 45 includes the computer-readable medium of any one of examples 38 to 44, including or excluding optional features. In this example, the plurality of views include multiple views of a calibration target.

[0104] Example 46 includes the computer-readable medium of any one of examples 38 to 45, including or excluding optional features. In this example, calibration is performed without prior knowledge of a calibration target.

[0105] Example 47 includes the computer-readable medium of any one of examples 38 to 46, including or excluding optional features. In this example, averaging the three dimensional data and applying a subpixel refinement to the plurality of features results in a smoothing of the depth data.

[0106] It is to be understood that specifics in the aforementioned examples may be used anywhere in one or more embodiments. For instance, all optional features of the computing device described above may also be implemented with respect to either of the methods or the computer-readable medium described herein. Furthermore, although flow diagrams and/or state diagrams may have been used herein to describe embodiments, the techniques are not limited to those diagrams or to corresponding descriptions herein. For example, flow need not move through each illustrated box or state or in exactly the same order as illustrated and described herein.

[0107] The present techniques are not restricted to the particular details listed herein. Indeed, those skilled in the art having the benefit of this disclosure will appreciate that many other variations from the foregoing description and drawings may be made within the scope of the present techniques. Accordingly, it is the following claims including any amendments thereto that define the scope of the present techniques.

CLAIMS

What is claimed is:

1. An apparatus for calibrating multiple cameras, comprising:
an image capture module to capture depth data and sensor data for a plurality of views;
an extraction module to extract a first plurality of features from the depth data and a second plurality of features from the sensor data for each view;
a correspondence module to locate corresponding features in the first plurality of features and the second plurality of features for each view;
a depth module to generate three-dimensional data for each feature of the first plurality of features for each view; and
a calibration module to calibrate the multiple cameras by matching the generated three dimensional data with the corresponding features in the first plurality of features and the second plurality of features.
2. The apparatus of claim 1, wherein a controller is to average the three-dimensional data for each view, and the calibration module is to calibrate the multiple cameras by matching the averaged three dimensional data with the corresponding features in the first plurality of features and the second plurality of features.
3. The apparatus of claim 1, comprising averaging the three dimensional data across multiple views by calculating a best fit rigid transform from each view of a plurality of views to a reference view.
4. The apparatus of claim 1, comprising averaging the three dimensional data by tracking a camera pose as each view of the plurality of views is captured and using the camera pose to average the views into a single view.

5. The apparatus of claim 1, wherein the depth data is obtained from a stereoscopic infrared image pair.
6. A method for calibrating multiple cameras, comprising:
capturing depth data and sensor data for a plurality of views;
extracting a plurality of corresponding features from the depth data and the sensor data;
generating three-dimensional data for each feature of the plurality of corresponding features; and
calibrating the multiple cameras by calculating a projection based on the three-dimensional data and the plurality of corresponding features.
7. The method of claim 6, wherein the three-dimensional data is averaged across the multiple views, and the multiple cameras are calibrated by calculating a projection based on the averaged three-dimensional data and the plurality of corresponding features across multiple views.
8. The method of claim 6, comprising averaging the three dimensional data across multiple views by calculating a best fit rigid transform from each view of a plurality of views to a reference view.
9. The method of claim 6, comprising averaging the three dimensional data by tracking a camera pose as each view of the plurality of views is captured and using the camera pose to average the views into a single view.
10. The method of claim 6, wherein the depth data is obtained from a stereoscopic infrared image pair.
11. A system for calibrating multiple cameras, comprising:
a depth camera and a sensor to capture images;
a memory configured to receive image data; and
a processor coupled to the memory, depth camera, and sensor, the processor to:

capture depth data and sensor data for a plurality of views;
extract a plurality of corresponding features from the depth data and
the sensor data;
generate three-dimensional data for each feature of the plurality of
corresponding features; and
calibrate the multiple cameras by calculating a projection based on the
three-dimensional data and the plurality of corresponding
features.

12. The system of claim 11, wherein the processor is to average the three-dimensional data for each view, and the calibration module is to calibrate the multiple cameras by matching the averaged three dimensional data with the corresponding features in the first plurality of features and the second plurality of features.

13. The system of claim 11, comprising averaging the three dimensional data across multiple views by calculating a best fit rigid transform from each view of a plurality of views to a reference view.

14. The system of claim 11, comprising applying subpixel refinement to each feature of the plurality of features.

15. The system of claim 11, wherein features are detected by observing points of interest in each view of the plurality of views.

16. An apparatus for calibrating multiple cameras, comprising:
an image capture module to capture depth data and sensor data for a plurality of views;
a means to extract a first plurality of features from the depth data and a second plurality of features from the sensor data for each view;
a means to locate corresponding features in the first plurality of features and the second plurality of features for each view;
a means to generate three-dimensional data for each feature of the first plurality of features for each view; and

a means to calibrate the multiple cameras by matching the generated three dimensional data with the corresponding features in the first plurality of features and the second plurality of features.

17. The apparatus of claim 16, wherein a means to average views is to average the three-dimensional data for each view, and the means to calibrate the multiple cameras is to match the averaged three dimensional data with the corresponding features in the first plurality of features and the second plurality of features.

18. The apparatus of claim 16, comprising averaging the three dimensional data across multiple views by calculating a best fit rigid transform from each view of a plurality of views to a reference view.

19. The apparatus of claim 16, comprising averaging the three dimensional data by tracking a camera pose as each view of the plurality of views is captured and using the camera pose to average the views into a single view.

20. The apparatus of claim 16, wherein the depth data is obtained from a stereoscopic infrared image pair.

21. At least one machine readable medium comprising a plurality of instructions that, in response to being executed on a computing device, cause the computing device to:

capture depth data and sensor data for a plurality of views;

extract a plurality of corresponding features from the depth data and the sensor data;

generate three-dimensional data for each feature of the plurality of corresponding features; and

calibrate the multiple cameras by calculating a projection based on the three-dimensional data and the plurality of corresponding features.

22. The computer readable medium of claim 21, wherein the depth data is obtained from a stereoscopic infrared image pair.

23. The computer readable medium of claim 21, wherein the sensor data is an RGB sensor data.

24. The computer readable medium of claim 21, wherein the depth data is obtained from structured light.

25. The computer readable medium of claim 21, wherein calibration is performed without prior knowledge of a calibration target.

1/8






FIG. 1

2/8




200A
FIG. 2A

3/8




200B
FIG. 2B

4/8




300
FIG. 3

5/8



400
FIG. 4A


6/8



400B

FIG. 4B

7/8



500

FIG. 5

8/8




FIG. 6

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US2017/014825**A. CLASSIFICATION OF SUBJECT MATTER****H04N 13/02(2006.01)i, G06T 7/80(2017.01)i**

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)

H04N 13/02; H04N 13/00; G06T 3/40; G06T 17/20; G06T 7/00; H04N 17/00; G06T 7/80

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

Korean utility models and applications for utility models

Japanese utility models and applications for utility models

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

eKOMPASS(KIPO internal) & keywords: camera calibration, depth, sensor, correspondence

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	US 2016-0012588 A1 (MITSUBISHI ELECTRIC RESEARCH LABORATORIES, INC. et al.) 14 January 2016 See paragraphs [0012]-[0028]; and figures 1-3.	1-25
Y	US 2016-0073102 A1 (INTEL CORPORATION) 10 March 2016 See paragraphs [0011]-[0026]; claim 1; and figures 1-2.	1-25
A	US 2015-0213589 A1 (ALTEK SEMICONDUCTOR CORP.) 30 July 2015 See paragraphs [0025]-[0039]; and figures 1-3B.	1-25
A	US 2015-0213647 A1 (INTERMEC IP CORP.) 30 July 2015 See paragraphs [0074]-[0086]; and figures 3-4.	1-25
A	US 2013-0010079 A1 (CHA ZHANG et al.) 10 January 2013 See paragraphs [0044]-[0048]; and figures 3-4.	1-25



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

25 April 2017 (25.04.2017)

Date of mailing of the international search report

25 April 2017 (25.04.2017)

Name and mailing address of the ISA/KR

International Application Division

Korean Intellectual Property Office

189 Cheongsa-ro, Seo-gu, Daejeon, 35208, Republic of Korea

Facsimile No. +82-42-481-8578

Authorized officer

AHN, Jeong Hwan

Telephone No. +82-42-481-8633



INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No.

PCT/US2017/014825

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2016-0012588 A1	14/01/2016	WO 2016-009811 A1	21/01/2016
US 2016-0073102 A1	10/03/2016	TW 201622419 A	16/06/2016
		US 9491452 B2	08/11/2016
		WO 2016-036470 A1	10/03/2016
US 2015-0213589 A1	30/07/2015	TW 201530493 A	01/08/2015
		TW I511081 B	01/12/2015
US 2015-0213647 A1	30/07/2015	US 2013-293540 A1	07/11/2013
		US 9007368 B2	14/04/2015
		US 9292969 B2	22/03/2016
US 2013-0010079 A1	10/01/2013	US 9270974 B2	23/02/2016
		WO 2013-009662 A2	17/01/2013
		WO 2013-009662 A3	07/03/2013