

(21) Application No: 1520953.9

(22) Date of Filing: 27.11.2015

(71) Applicant(s):
Fuel 3D Technologies Limited
(Incorporated in the United Kingdom)
Unit 2 Douglas Court, Seymour Business Park,
Station Road, Chinnor, Oxfordshire, OX39 4HA,
United Kingdom

(72) Inventor(s):
Andrew Henry John Larkins
Willaim Richard Dendy

(74) Agent and/or Address for Service:
Marks & Clerk LLP
Fletcher House (2nd Floor), Heatley Road,
The Oxford Science Park, OXFORD, OX4 4GE,
United Kingdom

(51) INT CL:
G06T 7/55 (2017.01) **G06T 17/00** (2006.01)

(56) Documents Cited:
GB 2458927 A **GB 2453163 A**
US 20140078296 A1 **US 20110285822 A1**

(58) Field of Search:
INT CL **G01J, G02B, G06T, H04N**
Other: **WPI, EPODOC, X-FULL - designating the full text databases of all states**

(54) Title of the Invention: **Systems and methods for forming three-dimensional models of objects**
Abstract Title: **Creating 3D models from images captured at different light intensities**

(57) Creating a 3D model of an object comprising: illuminating the object a plurality of times at different respective intensities from at least three directions using at least one directional energy source (eg. LED or flash light source); capturing images of the object for each direction at each time; obtaining photometric data from the images; generating the 3D model using the photometric data. An energy intensity ratio may be determined. An ambient light photograph may be subtracted from the collected pictures. An initial three dimensional model may be created (eg. by stereoscopy) and then refined (eg. by determining the normal direction to the object's surface). Brightness values for points on the object may be determined and combined. Brightness and intensity of a range of colours may be acquired. Distance between the light source and points on the object may be compensated by weighting. Dynamic range of the frames may be reduced by a tone mapping algorithm.

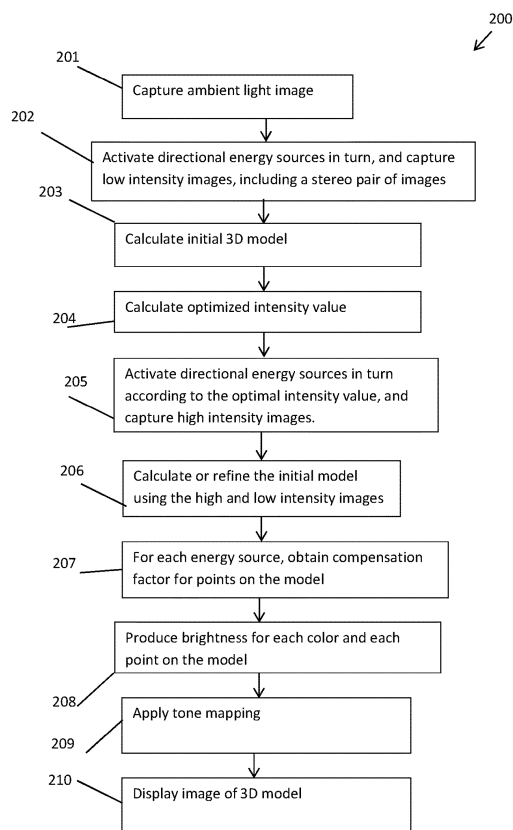


Fig. 3




Fig. 1




Fig. 2

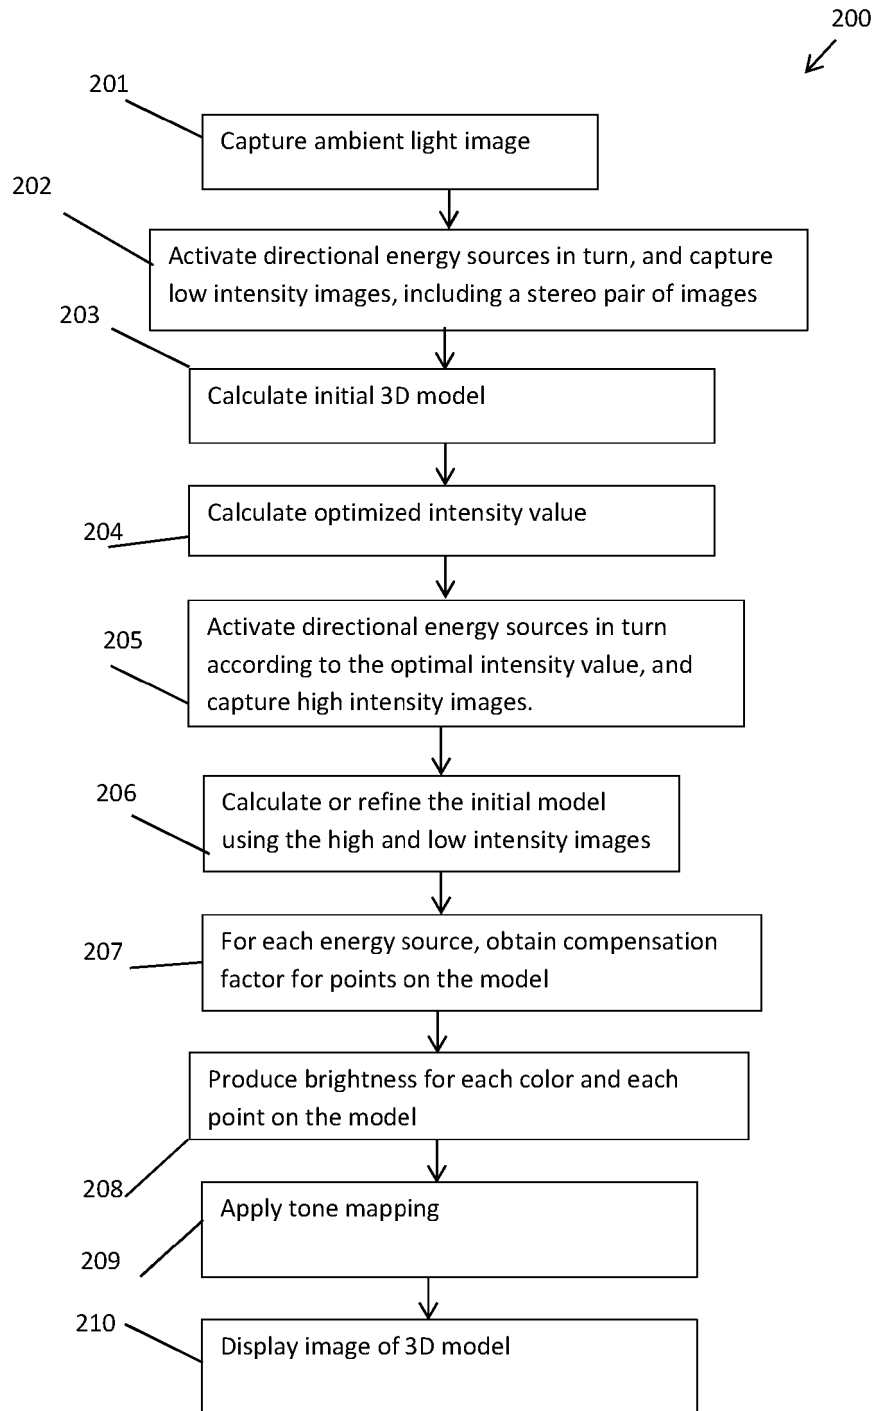


Fig. 3

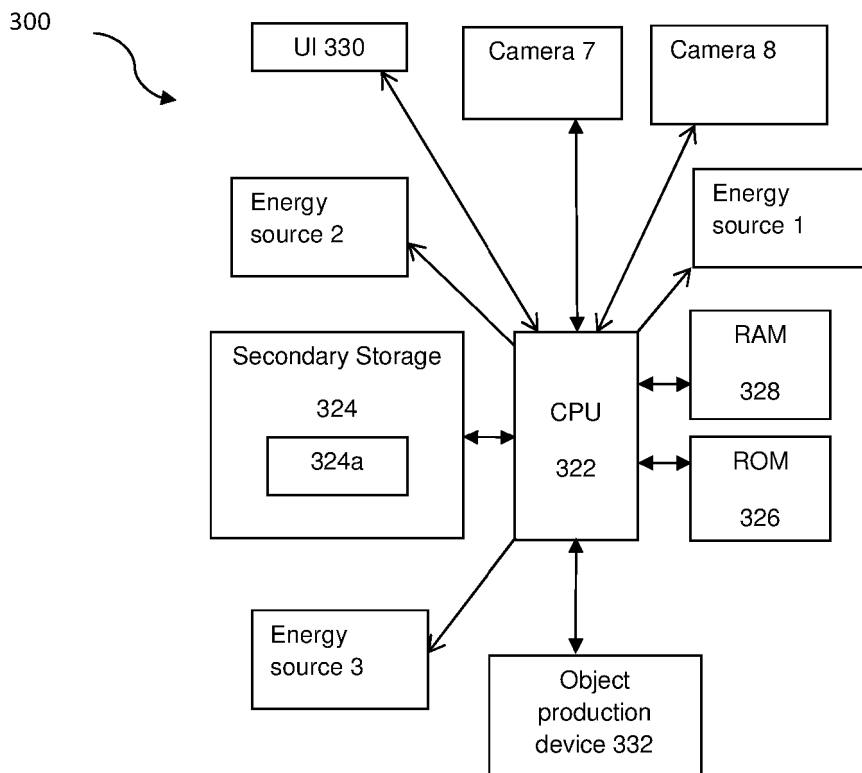


FIG. 4



The following terms are registered trade marks and should be read as such wherever they occur in this document:

ON Semiconductor (Page 9)

Systems and methods for forming three-dimensional models of objects

Summary of the invention

The present invention relates to an imaging system for generating three-dimensional (3D) models of a 3D object from two dimensional images, and a method performed by the imaging system.

Background of the invention

Modelling of 3D surfaces using two-dimensional images has been a major research topic for many years. The 3D surface is illuminated by light (or other electromagnetic radiation), and the two-dimensional images are created using the light reflected from it.

Most real objects exhibit two forms of reflectance: specular reflection (particularly exhibited by glass or polished metal) in which, if incident light (visible light or other electromagnetic radiation) strikes the surface of the object in a single direction, the reflected radiation propagates in a very narrow range of angles; and Lambertian reflection (exhibited by diffuse surfaces, such as matte white paint) in which the reflected radiation is isotropic with an intensity according to Lambert's cosine law (an intensity directly proportional to the cosine of the angle between the direction of the incident light and the surface normal). Most real objects have some mixture of Lambertian and specular reflective properties.

Recently, great progress has been made in imaging three-dimensional surfaces which exhibit Lambertian reflective properties by means of photometry (the science of measuring the brightness of light). For example, WO 2009/122200, "3D Imaging System" describes a system in which, in preferred embodiments, the object is successively illuminated by at least three directional light sources, and multiple cameras at spatially separated positions capture images of the object. If the object is expected to move relative to the imaging system during the imaging process, a localization template, fixed to the object, is provided in the optical fields of all the light sensors, to allow the images to be registered with each other, in a frame of reference in which the object is unmoving.

Typically, the object will have a number of "landmarks" which, when imaged, produce features which can be recognized in each of the images. Consider two images (a "stereo pair" of images) which are captured simultaneously respectively by two or more of the cameras which have a

known geometric relationship. For each of a number of landmarks (such as marks on the surface, or places where there is a color transition), the system determines the corresponding positions in the stereo pair of images of the corresponding features. Using this data, an initial 3D model of the object is created stereoscopically (i.e. by optical triangulation).

5 Photometric data is generated from images captured at different times when successive ones of the directional light sources are activated. If the object is moving relative to the cameras during this period, the images are registered using the localization template (i.e. such that the respective viewpoints are known in a common reference frame in which the object is stationary). On the assumption that the object exhibits Lambertian reflection, the photometric data makes it
10 possible to obtain an estimate of the normal direction to the surface of the object with a resolution comparable to individual pixels of the image. The normal directions are then used to refine the initial model of the 3D object.

Problems may arise due to the “near far” effect, in that areas of the object which are further from the light sources are lit less well than areas which are closer. This leads to lighting with a high
15 dynamic range. In any area of the object which is relatively darker, the variations of intensity due to variations in the normal direction are correspondingly lower, leading to less accurate results. For example, if a given captured image contains both bright areas and dark areas, and if the intensity of any given pixel is expressed as, say, an 8-bit number, then the small absolute difference in intensity between different pixels in the darker areas of the image, will be
20 expressed inaccurately. This limits the useful range of the photometric system.

A further problem is that typically the object is illuminated not only by the directional energy sources but by ambient light, which propagates in a direction which is not known and with unknown intensity. Thus, the light collected when one of the light sources is activated is subject to a background noise. This reduces the accuracy of the photometry. The noise due to the
25 ambient light can be reduced by making the light sources much brighter than the ambient light, but this can lead to areas of the collected images which are over-exposed because the object brightly reflects the light from the light sources. Over-exposure means that the intensity value of each pixel of the image does not accurately express the intensity of reflected light, leading to further problems for the photometry.

30 Summary of the invention

The present invention aims to provide new and useful methods and systems for obtaining three-dimensional (3D) models of a 3D object, and displaying images of the models.

In general terms the invention proposes a 3D modelling system in which an object is illuminated by at least one directional energy source (typically successively) from at least three directions.

5 The directional energy source(s) are controlled to emit energy at different respective intensities at different times. For each direction, at least one image of the object is captured at each of multiple corresponding light intensities. The images captured when the object is illuminated in one of the directions by the at least one energy source are here referred to as "captured images".

10 The multiple images of the object for each direction provide high dynamic range data about the propensity of points on the object to reflect light propagating in the corresponding direction.

Note that the problem of imaging a visual field which has wide variation in brightness (i.e. high dynamic range) is not unique to the field of 3D modeling using photometry. It appears also in the field of photographing scenes with an HDR, to produce visually pleasing two-dimensional
15 images. In that field, the conventional way to address HDR is tone mapping, in which the camera converts light intensities in the scene into pixel intensities in the image using a non-linear mapping. Thus, dark areas of the scene with relatively low levels of variation in absolute intensity appear in the image as regions with a higher range of pixel intensities. Some camera systems include tone mapping as a standard feature. However, the present inventors have
20 realized that using tone mapped data for photometric 3D modeling of an object would result in distortions in the reconstructed 3D model, because the pixel intensities would not be proportional to the intensity of light reflected from the object.

Another conventional way of addressing HDR is by capturing images of the visual field with two or more different exposure times: a long exposure image in which the areas corresponding to
25 the brightest parts of the visual field are over-exposed; and a short exposure image in which the areas corresponding to the darkest parts of the visual field are almost black. The two images are merged into a single image by tone mapping, so that areas of the merged image corresponding to dark areas of the visual field tend to be similar to the long exposure image, while areas of the merged image corresponding to bright areas of the visual field tend to be
30 similar to the short exposure image. However, this too leads to a non-linear relationship between brightness in the visual field and brightness in the merged image, which leads to distortions in the 3D modelling based on photometry. Furthermore, the long exposure image will

include more ambient light, which also leads to a reduction in the quality of the 3D model. By contrast, in the present invention, all images for a given directional energy source may be captured with a substantially similar exposure time, so that ambient light affects all equally.

Note furthermore that since the areas of the object which are brightest will be different according to the direction from which the object is illuminated, conventional techniques for dealing with HDR will operate differently for each direction. That is, the tone mapping function would be different for different ones of the directions. This inconsistency would lead to additional errors when the images for the different directions are used together for the 3D modeling by photometry.

By contrast, the present invention makes possible imaging in which, for a given illumination direction, accurate reflectivity values are obtained for both the bright and dark areas of the object. Preferably two images are captured for each direction. Preferably no tone mapping is applied to either image: in each image the intensity level of each pixel is directly proportional to the brightness of the corresponding area of the object (metrically accurate luminance).

Furthermore, the images for different ones of the directions can be obtained in a consistent way. For example, preferably, the object is illuminated from each direction with the same ratio of its respective brightness levels when the darkest image and brightest image are captured.

Preferably, for each direction, a respective image is captured when the energy source has a relatively low intensity (a "low intensity image"). The low intensity images for the various directions are used to obtain an energy intensity ratio. A respective second image is then captured for each direction using a higher "optimized" intensity, which is higher than the low intensity by the energy intensity ratio. This allows the optimized intensity value to be chosen to make a preferred trade-off between the need for the intensity value to be high enough to sufficiently illuminate the darker areas of the object, and the desirability of avoiding over-exposure.

An appropriate selection of the higher intensity value makes it possible to minimize the total number of images which have to be taken for each of the directions to ensure that all areas of the object are imaged in at least one of the captured images with a pixel intensity which is an accurate expression of their brightness. Reducing the number of images captured for each direction reduces the amount of data processing required to analyze them. Furthermore, it reduces the time taken to capture the images, and thus the risk that the object moves relative to

the imaging system during the imaging. From another point of view, it makes it easier to correct the images for relative motion of the object and imaging system during the imaging process.

Preferably, in addition to images captured when the directional energy source(s) are illuminated, the imaging system captures at least one image when the directional energy source(s) are not illuminated and the object is lit only by ambient light. The intensity of this "ambient light image" can be used to correct the images captured when the directional energy source(s) are illuminated, thereby removing the effect of the ambient light. This can be done simply by subtracting the ambient light image from the captured images, since the captured images have pixel intensities which are directly proportional to the intensity of the light transmitted by the corresponding areas of the object.

Preferably, the photometric information obtained from the images captured when the object is illuminated from the successive directions, is used to provide fine detail to improve an initial 3D model obtained in another way. As in WO 2009/122200, the initial 3D model may be obtained by stereoscopy, using at least a stereo pair of images of the object captured from different viewpoints (for example simultaneously). Alternatively, it may be obtained by a depth camera. Known types of depth camera include those using sheet-of-light triangulation, structured light (that is, light having a specially designed light pattern), time-of-flight or interferometry.

As mentioned above, the object is preferably illuminated successively by each energy source in at least three directions. This is done with the low intensity first, and then with the high intensity. If this is done, the energy sources may emit light of the same frequency spectrum (e.g. if the energy is visible light, the directional light sources may each emit white light, and the captured images may be color images). However, in principle, for either or both of the intensity levels, the object could alternatively be illuminated in at least three directions by energy sources which emit energy with different respective frequency spectra (e.g. in the case of visible light, the directional light sources may respectively emit red, white and blue light). In this case, the directional energy sources could be activated simultaneously (once at the low intensity and once at the high intensity), if the energy sensors are able to distinguish the energy spectra. For example, the energy sensors might be adapted to record received red, green and blue light separately. That is, the red, green and blue light channels of the captured images would be captured simultaneously, and would respectively constitute the images in which the object is illuminated in a single direction. However, this second possibility is not preferred, because coloration of the object may lead to incorrect photometric imaging.

Although photometric imaging requires that the object is illuminated from at least three directions, the number of directions may be higher than this. The timing of the activation of the energy source(s) and energy sensor(s) may be controlled by a processor, which may also calculate relative motion of the object and energy sensor(s), and compensate for it.

- 5 In principle, there could be only one directional energy source which is moved to as to successively illuminate the object from successive directions.

However, typically at least three energy sources are provided, for illuminating the object from respective directions. It would be possible for these sources to be provided as at least three energy outlets from an illumination system in which there are fewer than three elements which
10 generate the energy. For example, there could be a single energy generation unit (light generating unit) and a switching unit which successively transmits energy generated by the single energy generation unit to respective input ends of at least three energy transmission channels (e.g. optical fibers). The energy would be output at the other ends of the energy transmission channels, which would be at three respective spatially separate locations. Thus
15 the output ends of the energy transmission channels would constitute respective energy sources. The light would propagate from the energy sources in different respective directions.

The at least three directional energy sources may be provided in a plane perpendicular to a line of sight of one of the energy sensors, and symmetrically arranged around that line of sight. Optionally, an additional directional energy source may provided close to at least one of the
20 energy sensor(s). This provides “bright field” lighting, i.e. ensuring that the whole of the object which is visible to the at least one energy sensor is lit to some extent, so that there are no completely dark regions.

Once constructed, the 3d model of the object may be displayed using a screen, which typically shows the 3D model from a certain perspective, with certain coloring and texture. The coloring
25 and texture may be derived from at least some of the captured images (the images taken when the object is illuminated by one of the directional energy sources). This may include compensating intensity in the captured images to take into account the distance of any point in the model from the corresponding directional energy source (i.e. compensating for the near-far effect). Thus, the intensity of a pixel in the displayed image of the model corresponding to a
30 point of the model may be represented with a brightness which is derived using the images captured by the directional energy sources but compensated by a compensation factor which is obtained by: calculating the distance of the point in the model from the energy source(s), and

calculating the compensation factor based on the distance (for example, the compensation factor may be the square of the distance, which compensates for the usual inverse square diminution of brightness with distance).

5 If desired, the dynamic range of the intensity of pixels in the representation of the 3D model may be reduced by one of the techniques explained above in relation to two-dimensional imaging, such as tone mapping.

10 The energy used is electromagnetic radiation, i.e. light. The term "light" is used in this document to include electromagnetic radiation which is not in the visible spectrum. Various forms of directional energy source may be used in embodiments of the invention. For example, a standard photographic flash, a high brightness LED cluster, or Xenon flash bulb or a 'ring flash'. It will be appreciated that the energy need not be in the visible light spectrum. One or more of the energy sources may be configured to generate light in the infrared (IR) spectrum (wavelengths from 900nm to 1mm) or part of the near infrared spectrum (wavelengths from 900nm to 1100nm). Optionally, the energy may be polarized.

15 Where visible-light directional energy is applied, then the energy sensors may be two or more standard digital cameras, or video cameras, or CMOS sensors and lenses appropriately mounted. In the case of other types of directional energy, sensors appropriate for the directional energy used are adopted. A discrete energy sensor may be placed at each viewpoint, or in another alternative a single sensor may be located behind a split lens or in combination with a
20 mirror arrangement.

The energy sources and energy sensors preferably have a known positional relationship, which is typically fixed. The energy sensor(s) and energy sources may be incorporated in a portable apparatus, such as a hand-held instrument. Alternatively, the energy sensor(s) and energy sources may be incorporated in an apparatus which is mounted in a building.

25 The invention may be expressed as an apparatus for capturing images, including a processor for analyzing the images according to program instructions (which may be stored in non-transitory form on a tangible data storage device). Alternatively, it may be expressed as the method carried out by the apparatus.

Brief description of the drawings

An embodiment of the invention will now be described for the sake of example only with reference to the following figures in which:

Fig. 1 shows a first schematic view of an imaging assembly for use in an embodiment of the present invention to form a 3D model of an object, including multiple image capturing devices;

Fig. 2 is a view of the object from the opposite direction from Fig. 1;

Fig. 3 is a flow diagram of a method which is an embodiment of the invention; and

Fig. 4 illustrates an embodiment of the invention incorporating the imaging assembly of Fig. 1 and a processor.

Detailed description of the embodiments

Referring firstly to Fig. 1, an imaging assembly is shown which is a portion of an embodiment of the invention. The imaging assembly includes directional energy sources 1, 2, 3. It further includes energy sensors 7, 8 in form of image capturing devices. The image capturing devices capture color images, i.e. intensity values at each pixel for each of three colors. The energy sources 2, 3 are fixedly mounted to each other by a strut 6. The image capturing device 7 is fixedly mounted to the energy source 1 by the strut 4a, to the energy source 2 by strut 4b, and to the image capturing device 8 by the strut 12. The exact form of the mechanical connection between the energy sources 1, 2, 3 is different in other forms of the invention, but it is preferable if it maintains the energy sources 1, 2, 3 and the imaging devices 7, 8 at fixed distances from each other and at fixed relative orientations. Preferably, the energy sources 1, 2, 3 are symmetrical about (i.e. at equal angular positions around) the imaging device 8. The energy sources 1, 2, 3 may be substantially co-planar, in a plane perpendicular to the imaging direction of the imaging device 8. The energy sources 1, 2, 3 and image capturing devices 7, 8 may be incorporated in a portable, hand-held instrument. In addition to the imaging assembly shown in Fig. 1, the embodiment includes a processor 322 (not shown in Fig. 1) which is in electronic communication with the energy sources 1, 2, 3 and image capturing devices 7, 8. This is described below in detail with reference to Fig. 4.

The energy sources 1, 2, 3 are each adapted to generate electromagnetic radiation ("light"), such as visible light or infra-red radiation. The energy sources 1, 2, 3 and image capturing devices 7, 8, are all controlled by the processor. The output of the image capturing devices 7, 8

is transmitted to the processor. The energy sources 1, 2, 3 are powered by a power supply (not shown) which may include a capacitor for accumulating charge during the time before the energy sources are activated.

Each of the image capturing devices 7, 8 is arranged to capture an image of an object 9
5 positioned in both the respective fields of view of the image capturing devices 7, 8. The image capturing devices 7, 8 are spatially separated, and preferably also arranged with converging fields of view, so the apparatus is capable of capturing two images of the object 9 from two separated viewpoints (a "stereo pair" of images), so that stereoscopic imaging of the object 9 is possible. Note that the images captured are typically color images, having a separate intensity
10 for each pixel each of three color channels. In this case, the three channels are treated separately in the process described below. Alternatively, in variations of the embodiment, the three color channels could be combined together into a single channel (i.e. by at each pixel summing the intensities of the channels), or, in principle, two of the channels could be discarded.

15 The case of two viewpoints is often referred to as a "stereo pair" of images, although it will be appreciated that in variations of the embodiment more than two spatially-separated image capturing devices may be provided, so that the object 9 is imaged from more than two viewpoints. This may increase the precision and/or visible range of the apparatus. The words "stereo" and "stereoscopic" as used herein are intend to encompass, in addition to the possibility
20 of the subject being imaged from two viewpoints, the possibility of the subject being imaged from more than two viewpoints.

Suitable image capture devices for use in the invention include the 1/3-Inch CMOS Digital Image Sensor (AR0330) provided by ON Semiconductor of Arizona, US.

Fig. 2 shows schematically how the object 10 looks as viewed which is opposite to the viewing
25 direction of Fig. 1. This is the direction in which the image capturing devices 7, 8 face. As shown in both Figs. 1 and 2, a localization template 10 is present in the visual field of both the image capturing devices 7, 8, and in a substantially fixed positional relationship with the object 9. The localization template 10 is useful, though not essential, for registering the images in relation to each other, particularly in the case that the object 9 may move relative to the image capturing
30 devices 7,8 during the imaging process. Since it is in the visual field of both the image capturing devices 7, 8, it appears in all the images captured by those devices, and it is provided with a known pattern, so that the processor is able to identify it from the image, and from its position,

size and orientation in any given one of the images, reference that image to a coordinate system defined in relation to the localization template 10. In this way, all images captured by the image capturing devices 7, 8 can be referenced to that coordinate system. If the object 9 moves slightly between the respective times at which any two successive images are captured, the localization template 10 will move correspondingly, so the object 9 will not have moved in the coordinate system. In variations of the embodiment in which the positional relationship of the energy sources 1, 2, 3 and image capturing devices 7, 8 is not known, it may be determined if the energy sources 1, 2, 3 illuminate the localization template 10. Note that if the object 9 is known to be stationary relative to the imaging system, the localization template 10 may be unnecessary. In other embodiments of the invention in which the object 9 may move relative to the imaging system, the images captured by image capturing devices 7, 8 may be mutually registered in other ways, such as identifying in each image landmarks of the object 9, and using these landmarks to register the images with each other.

Turning to Fig. 3 a method 200 is shown which employs the imaging assembly of Fig. 1. In step 201 of method 200 the processor controls the image capturing device 8 capture an image of the object 9 at a time when the energy sources 1, 2, 3 are all turned off (an "ambient light image").

In step 202, each of the energy sources 1, 2, 3 is successively activated to generate energy. This is done with a first energy intensity. Supposing that the power supply has accumulated an energy denoted by $3E$ (that is, energy E for each of the energy sources 1, 2, 3), the first energy intensity is chosen such that less than half of the energy is expended. For example, the first energy intensity may be chose such that about 20% of energy E is supplied to each of the energy sources.

While each of the energy sources 1, 2, 3 is activated successively, the image capturing device 8 captures three respective images ("low intensity images"). While one of the low intensity images is being captured, the image capturing device 7 also captures an image, so that a stereo pair of simultaneously captured images is created. Note that in a variation of the embodiment, step 202 can be performed before step 201, or step 201 can be performed mid-way through step 202, i.e. when some but not all of the lower intensity images have been captured.

In optional step 203, the processor uses the stereo pair of images geometrically, e.g. by the same stereoscopic algorithm employed in WO 2009/122200, to produce an initial 3D model of the object 9. This is based around known principles of optical parallax. This technique generally provides good unbiased low-frequency information (the coarse underlying shape of the surface

of the object), but is noisy or lacks high frequency detail. The stereoscopic reconstruction uses optical triangulation, by geometrically correlating pairs of features in the respective stereo pair of images captured by the image capture devices 7, 8 to give the positions of each of the corresponding landmarks on the object 9 in a three-dimensional space defined based on the localization template 10. Note that this step is computationally expensive, and therefore time-consuming unless the system includes a fast processor, so certain embodiments may omit it in the interests of keeping steps 202 and 205 as close to each other in time as possible. Other embodiments may perform step 203 but forming a low-resolution 3D model, thereby keeping the computational cost low.

10 In step 204 the processor 322 uses the low intensity images to calculate an optimized intensity value for capturing second images for each of the energy sources 1, 2, 3 (in step 205). Each of the energy sources has the same ratio of the optimized intensity value and the low intensity used in step 202. This ratio is referred to as the “energy intensity ratio”. In fact, it may be that the low intensity used in step 202 is the same for all energy sources 1, 2, 3, in which case the
15 optimized intensity value is the same for all the energy sources 1, 2, 3 also.

There are many ways in which the energy intensity ratio can be obtained within the scope of the invention.

In one example, initially the processor 322 may form a combined image from the low intensity images. For example, each pixel of the combined image may have an intensity for each color which is the mean of the respective intensities for the corresponding color of the corresponding
20 pixels of the three low intensity images. Or the combined image may have an intensity for each color which is the highest of the respective intensities for the corresponding color of the corresponding pixels of the three low intensity images. The processor 322 then combines the intensities for each color to obtain a single combined intensity value for each pixel. If optional
25 step 203 was performed, the initial 3D model is used to identify pixels which are images of the background (i.e. not parts of the object 9), and such pixels are removed from consideration. For some of the remaining pixels, the combined intensity value is above a threshold, indicating that the corresponding points on the object (“bright areas”) were adequately imaged in the low intensity images. These pixels are excluded, and the processor 322 then forms an average
30 (“dark area intensity”) of the combined intensities of the remaining pixels (corresponding to “dark areas” of the object). The processor 322 forms the optimized intensity value as a predefined function of the dark area intensity. The optimized intensity value is higher if the dark area

intensity is lower. Thus, a low value for the dark area intensity may indicate that it is worth risking a higher value for the optimized intensity value, even though this risks over-exposure in the bright areas of the object, in order to ensure that darker areas of the object are properly imaged.

5 The scheme above may have a number of variations. Firstly, since some of the dark area pixels may actually be showing the background of the object (i.e. not the object itself), pixels of the low energy images which have a combined intensity below a second threshold may be excluded from the calculation of the dark area intensity. Second, the intensity of the bright area pixels may also be used in the calculation of the optimized intensity value, e.g. such that if the average
10 intensity of the bright area pixels is higher, then optimized intensity value is chosen to be lower, to reduce over-exposure in the bright areas.

An alternative method for calculating the energy intensity ratio is based on the fact that the operator of the imaging system is likely to place the most important portion of the object 9 in the centre of the visual fields of the image capturing devices 7, 8. Accordingly, the processor 322
15 forms the optimized intensity value as a function of the respective intensities of the pixels in a central region of the visual fields of the image capturing devices 7, 8 ("central pixels"). The intensities input to the function may be the respective intensities of the central pixels in the low intensity images, or the respective intensities of the pixels in the combined image mentioned above. The function is chosen to try to ensure that in the second images the brightness of as
20 many of these pixels as possible will be within a predefined optimal brightness range, which ensures that the corresponding portion of the object is correctly imaged.

Note that in a variation of the embodiment, step 204 is performed before step 203, or the two may be performed concurrently.

In step 205, the processor 322 controls the energy sources 1, 2, 3 to successively generate
25 energy with an intensity according to the calculated optimized intensity value. While this is done, the image capturing device 8 captures three respective images ("high intensity images"). Thus, the high intensity images and the low intensity images collectively provide reliable information about the reflectivity of both bright and dark areas of the object 9. Both the low intensity and high intensity images have intensity values at each pixel which are directly proportional to the
30 level of energy emitted from the corresponding portion of the object 9 when the images were taken.

In step 206, the processor refines the initial 3D model obtained in step 203 using the high intensity and low intensity images. Alternatively, in embodiments in which step 203 was omitted, or in which it was performed to form a 3D model with low resolution, in step 206 a 3D model is calculated without reference to a previously calculated 3D model.

- 5 In either case, if there is a possibility that the object 9 moved relative to the imaging system in the time between the performance of steps 201 and 205, motion compensation may be applied to take register the ambient light image, high intensity images and low intensity images in a common reference frame, for example using the localization template 10.

10 The intensity of the ambient light image is subtracted from both the high intensity and low intensity images, thereby removing the effects of ambient light.

The rest of step 206 may be performed using the photometric technique employed in WO 2009/122200. The photometric reconstruction requires an approximating model of the surface material reflectivity properties. In the general case this may be modelled (at a single point on the surface) by the Bidirectional Reflectance Distribution Function (BRDF). A simplified model is typically used in order to render the problem tractable. One example is the Lambertian Cosine Law model. In this simple model the intensity of the surface as observed by the camera depends only on the quantity of incoming irradiant energy from the energy source and foreshortening effects due to surface geometry on the object. This may be expressed as:

$$I = P\rho \mathbf{L} \cdot \mathbf{N} \quad (\text{Eqn 1})$$

- 20 where I represents the intensity observed by the image capture device 8 at a single point on the object 9, P the incoming irradiant light energy at that point, \mathbf{N} the object-relative surface normal vector, \mathbf{L} the normalized object-relative direction of the incoming lighting and ρ the Lambertian reflectivity of the object at that point. Typically, variation in P and \mathbf{L} is pre-known from a prior calibration step, or from knowledge of the position of the energy sources 1, 2, 3, and this (plus the knowledge that \mathbf{N} is normalized) makes it possible to recover both \mathbf{N} and ρ at each pixel. Since there are three degrees of freedom (two for \mathbf{N} and one for ρ), intensity values I are needed for at least three directions \mathbf{L} in order to uniquely determine both \mathbf{N} and ρ . Note that typically only the single-source images captured by the image capturing device 8 are used in the photometry. This is why three energy sources 1, 2, 3 are provided. For a bright area pixel the processor 322 can use an intensity I which is calculated using the low intensity images, whereas for a dark area pixel the processor 322 can use an intensity I which is calculated using the high intensity images but taking

into account that the value of P is higher in the high intensity images in accordance with the optimized intensity value.

The shape of the 3D model of the object 9 (i.e. the surface of the object) has now been completed. The remaining steps of the method 200 relate to forming colors and brightness levels for the points on the model, and displaying the model.

In step 207, the processor 322 works out the distance of each point on the refined model from each of the energy sources. Then for each point, and for each of the energy sources, it works out a corresponding compensation factor, which may for example be the square of the distance.

In step 208, the processor 322 calculates a color and brightness for each point on the model. For a point in the bright area of the object it does this by, for each color, adding the respective intensities of that color in each of the low intensity images multiplied by the respective compensation factor. For a point in a dark area of the object it does this by, for each color, adding the respective intensities of that color in each of the low intensity images multiplied by the respective compensation factor, and dividing the result by the optimized intensity factor. The result is a set of brightness values for each color and each point which accurately express the brightness of the point with the full dynamic range of the object 9.

In step 209, tone mapping is then used to produce a modified intensity value for each color with a smaller dynamic range. For example, a non-linear function may applied to the values obtained in step 208 (tone mapping). For example, if the brightness values obtained in step 208 have 16-bit accuracy, step 209 may produce a modified intensity value as a 8-bit number for each point and each color. An alternative to tone mapping is the “exposure fusion” technique proposed by Mertens et al in 2007 (see Computer Graphics Forum, Volume 28, Issue 1, pages 161–171, March 2009)

In step 210 an image of the 3D model is displayed, with each point having a brightness for each color which is based on the modified intensity value obtained in step 209.

Fig. 4 is a block diagram showing a technical architecture of the overall system 300 for performing the method.

The technical architecture includes a processor 322 (which may be referred to as a central processor unit or CPU) that is in communication with the cameras 7, 8, for controlling when they capture images and receiving the images. The processor 322 is further in communication with, and able to control the energy sources 1, 2, 3.

The processor 322 is also in communication with memory devices including secondary storage 324 (such as disk drives or memory cards), read only memory (ROM) 326, random access memory (RAM) 328. The processor 322 may be implemented as one or more CPU chips.

5 The system 200 includes a user interface (UI) 330 for controlling the processor 322. The UI 330 may comprise a touch screen, keyboard, keypad or other known input device. If the UI 330 comprises a touch screen, the processor 322 is operative to generate an image on the touch screen. Alternatively, the system may include a separate screen (not shown) for displaying images under the control of the processor 322.

10 The system 200 optionally further includes a unit 332 for forming 3D objects designed by the processor 322, such as a replica of the object 9 according to the 3D model. For example the unit 332 may take the form of a 3D printer. Alternatively, the system 200 may include a network interface for transmitting instructions for production of the objects to an external production device.

15 The secondary storage 324 is typically comprised of a memory card or other storage device and is used for non-volatile storage of data and as an over-flow data storage device if RAM 328 is not large enough to hold all working data. Secondary storage 324 may be used to store programs which are loaded into RAM 328 when such programs are selected for execution.

20 In this embodiment, the secondary storage 324 has an order generation component 324a, comprising non-transitory instructions operative by the processor 322 to perform various operations of the method of the present disclosure. The ROM 326 is used to store instructions and perhaps data which are read during program execution. The secondary storage 324, the RAM 328, and/or the ROM 326 may be referred to in some contexts as computer readable storage media and/or non-transitory computer readable media.

25 The processor 322 executes instructions, codes, computer programs, scripts which it accesses from hard disk, floppy disk, optical disk (these various disk based systems may all be considered secondary storage 324), flash drive, ROM 326, RAM 328, or the network connectivity devices 332. While only one processor 322 is shown, multiple processors may be present. Thus, while instructions may be discussed as executed by a processor, the instructions may be executed simultaneously, serially, or otherwise executed by one or multiple processors.

30 In an alternative implementation of the invention, the method may be carried out by the image signal processor (ISP) associated with the camera chip of one of the image capturing devices.

In a further implementation of the invention, the method may be carried out by an integrated circuit such as a FPGA (field programmable gate array) for controlling the image capturing devices and light sources.

5 Whilst the foregoing description has described exemplary embodiments, it will be understood by those skilled in the art that many variations of the embodiment can be made within the scope of the attached claims.

Claims

1. An apparatus for computing a three-dimensional (3D) model of an object, comprising:
 - at least one directional energy source arranged to directionally illuminate the object from at least three directions;
 - 5 one or more energy sensors arranged to capture images of the object;
 - a processor arranged to receive images captured by the energy sensors, and control the energy sources and the at least one energy sensor;
 - the processor being arranged to control the at least one energy source and one or more energy sensors to capture, for each direction, a plurality of images at respective times at which
 - 10 the at least one energy source generates energy at different respective intensity levels;
 - to obtain photometric data using the captured images; and
 - to generate the 3D model using the photometric data.
2. An apparatus according to claim 1 in which the processor is arranged, for a certain energy intensity ratio, to control the at least one energy source to generate energy for each of
 - 15 the directions at a first respective said time with a respective first intensity level and at a second respective said time with a second respective intensity level,
 - the ratio of the first and second respective intensity levels being equal to the energy intensity ratio.
3. An apparatus according to claim 2 in which the processor is arranged to analyze first
 - 20 ones of the images captured at the first times, to obtain the energy intensity ratio.
4. An apparatus according to claim 3 in which the second intensity levels are higher than the corresponding first intensity levels.

5. An apparatus according to any preceding claim in which the processor is further arranged:

(i) to control at least one of the energy sensors to capture an ambient light image of the object at a time when the at least one energy source is not illuminating the object, and

5 (ii) to subtract the ambient light image from the captured images.

6. An apparatus according to any preceding claim in which the processor is arranged to generate the 3D model by:

generating an initial 3D model of the object; and

refining the initial 3D model using the photometric data.

10 7. An apparatus according to any preceding claim in which the processor is operative:

to generate, for a plurality of points on the model, at least one corresponding brightness value using the captured images; and

to display an image of the model with respective points on the model having a brightness dependent upon the at least one corresponding brightness value.

15 8. An apparatus according to claim 7 in which:

the energy sensors are operative to capture said images with respective intensity values at each pixel for each of a plurality of colors; and

for each of the points on the model, there is a corresponding said brightness value for each of the plurality of colors.

20 9. An apparatus according to claim 7 or 8 in which the processor is arranged to generate the brightness values by, for each said point on the model:

for some or all of the energy sources, obtaining a corresponding compensation factor indicative of the distance between the energy source and the point on the model; and

25 deriving the brightness level by combining intensity values of pixels of the captured images corresponding to the point on the model, the intensity values being weighted by the corresponding compensation factor.

10. An apparatus according to any of claims 7 to 9 in which the brightness of each point of the models is a modified intensity value obtained from the captured images using a dynamic range reduction algorithm.
11. An apparatus according to claim 10 in which the dynamic range reduction algorithm is tone mapping.
12. A method for computing a three-dimensional (3D) model of an object, comprising:
- (a) illuminating the object in at least three directions using at least one directional energy source, the or each energy source illuminating the object in each direction at a plurality of times at different respective intensity levels;
 - (b) at each of said times capturing an image of the object;
 - (c) obtaining photometric data using the captured images; and
 - (d) generating the 3D model using the photometric data.
13. A method according to claim 12 in which, for a certain energy intensity ratio, in step (a) for each of the directions the ratio of a first respective said intensity level and a second respective said intensity level is equal to the energy intensity ratio.
14. A method according to claim 13 further including analyzing a subset of the images to obtain the energy intensity ratio.
15. A method according to claim 14 in which the subset of the images includes corresponding images for each of the directions captured when the at least one energy source is generating energy at the respective first energy level, the second respective energy level being higher than the first said energy level.
16. A method according to any of claims 12 to 15 further including:
- capturing an ambient light image of the object at a time when the at least one energy source is not illuminating the object, and
 - subtracting the ambient light image from the images captured when the object is illuminated by the at least one energy source.
17. A method according to any of claims 12 to 16, in which said step (d) is performed by:

generating an initial 3D model of the object; and

refining the initial 3D model using the photometric data.

18. A method according to claim 17 in which the initial 3D model is formed by:

5 capturing images of the object from multiple respective viewpoints, and analyzing the images captured from multiple respective inputs stereoscopically.

19. A method according to any of claims 12 to 18 further comprising:

(e) generating, for a plurality of points on the model, at least one corresponding brightness value using the captured images; and

10 (f) displaying an image of the model with respective points on the model having a brightness dependent upon the at least one corresponding brightness value.

20. A method according to claim 19 in which:

said images include respective intensity values at each pixel for each of a plurality of colors; and

15 for each of the points on the model, there is a corresponding said brightness value for each of the plurality of colors.

21. A method according to claim 19 or 20 in which step (e) is performed by, for each said point on the model:

for some or all of the energy sources, obtaining a corresponding compensation factor indicative of the distance between the energy source and the point on the model; and

20 deriving the brightness level by combining intensity values of pixels of the captured images corresponding to the point on the model, the intensity values being weighted by the corresponding compensation factor.

22. A method according to any of claims 19 to 21 in which the brightness of each point of the models is a modified intensity value obtained from the captured images using a dynamic range reduction algorithm.
25

23. A method according to claim 22 in which the dynamic range reduction algorithm is tone mapping.



Application No: GB1520953.9

Examiner: Mr Charles Durden

Claims searched: 1-23

Date of search: 22 March 2016

Patents Act 1977: Search Report under Section 17

Documents considered to be relevant:

Category	Relevant to claims	Identity of document and passage or figure of particular relevance
X	1, 2, 4, 12	US 2011/0285822 A1 (DAI et al.) - see eg. paragraphs 10, 14, 15, 49, 51, 56-58, 90 and Figs. 2a & 2b
A	-	GB 2458927 A (EYKONA TECHNOLOGIES LTD) - see eg. page 2 lines 19-34
A	-	GB 2453163 A (BLAIR) - pages 2 and 4
A	-	US 2014/0078296 A1 (MAMIYA) - see eg. paragraphs 19, 28

Categories:

X	Document indicating lack of novelty or inventive step	A	Document indicating technological background and/or state of the art.
Y	Document indicating lack of inventive step if combined with one or more other documents of same category.	P	Document published on or after the declared priority date but before the filing date of this invention.
&	Member of the same patent family	E	Patent document published on or after, but with priority date earlier than, the filing date of this application.

Field of Search:

Search of GB, EP, WO & US patent documents classified in the following areas of the UKC^X :

--

Worldwide search of patent documents classified in the following areas of the IPC

G01J; G02B; G06T; H04N

The following online and other databases have been used in the preparation of this search report

WPI, EPODOC, X-FULL - designating the full text databases of all states

International Classification:

Subclass	Subgroup	Valid From
G06T	0007/00	01/01/2006
G06T	0017/00	01/01/2006