



(51) International Patent Classification:  
*G06T 7/00* (2017.01)

(21) International Application Number:  
PCT/EP2019/083428

(22) International Filing Date:  
03 December 2019 (03.12.2019)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:  
18213981.6 19 December 2018 (19.12.2018) EP

(71) Applicant: **NOKIA TECHNOLOGIES OY** [FI/FI];  
Karakaari 7, 02610 Espoo (FI).

(72) Inventors: **SUBRAMANIAN, Mahendran**; Brain and Behaviour Lab, Imperial College London, Prince Consort Road, London SW7 2AZ (GB). **FAISAL, Aldo**; Brain and Behaviour Lab, Imperial College London, Prince Consort Road, London SW7 2AZ (GB).

(74) Agent: **NOKIA EPO REPRESENTATIVES**; Nokia Technologies Oy, Karakaari 7, 02610 Espoo (FI).

(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, JO, 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).

**Published:**

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

(54) Title: MOVEMENT INDICATION



(57) Abstract: An apparatus, method and computer program is described comprising: receiving video data for a scene; determining a movement measurement for at least some of a plurality of subframes of the video data; weighting the movement measurements to generate a plurality of weighted movement measurements, wherein the weighting is dependent on the subframe; and generating a movement indication for the scene from a combination of some or all of the weighted movement measurements.

FIG. 2

20



## Movement Indication

### Field

This specification relates to generation of a movement indication for a scene based on video data of the scene.

### Background

Motion data can be extracted from video data. However, there remains a need for alternative arrangements for generating a movement indication for a scene.

### Summary

In a first aspect, this specification describes an apparatus comprising: means for receiving video data for a scene; means for determining a movement measurement for at least some of a plurality of subframes (e.g. subdivisions of a frame of video data) of the video data; means for weighting the movement measurements to generate a plurality of weighted movement measurements, wherein the weighting is dependent on the subframe; and means for generating a movement indication for the scene from a combination (such as a sum) of some or all of the weighted movement measurements.

The weighting may be a binary weighting. For example, by providing a binary weighting, some subframes may be ignored when generating a movement indication for a scene from a combination of weighted movement measurements. In some embodiments, the weighting may be non-binary.

The weighting for at least some of said subframes may be at least partially dependent on one or more of: historical movement data for that subframe; a location of the subframe within the scene; and visual characteristics of the scene.

In some embodiments, one or more of the weightings are updated on-the-fly. For example, as movement data is received, weightings may be adjusted. For example, as more movement data for a particular subframe is obtained, historical movement data that might be relevant to the weighting of that subframe may be updated, such that the weighting can be updated in-use or on-the-fly.

In some embodiments, a final generated movement indication may be generated offline, on the basis of movement measurement data that is obtained and weighted on-the-fly (which may, for example, save storage requirements and/or reduce data processing requirements).

The movement measurements for said subframes may comprise an indication of a degree of movement between successive instances of the respective subframes.

In some embodiments, determining the movement measurement for said subframes of the video data may comprise: means for determining a mean distance measurement for the  
5       respective subframes of the received video data; and means for determining, for one or more of the mean distance measurements, whether said mean distance measurement differs from a mean distance for a corresponding subframe of a preceding (e.g. immediately preceding) video data frame instance.

10       The means for determining the movement measurements for said subframes of the video data may determine that a movement has taken place if a mean distance measurement between successive subframes is greater than a threshold value. The threshold may be based on a determined standard deviation threshold (e.g. the threshold may be five standard  
15       deviations from the mean).

The video data may comprise depth data for a plurality of pixels of an image of the scene. Moreover, the movement measurements for said subframes of the video data may be determined on the basis of a change in received depth data for the respective subframe. The  
20       received depth data may, for example, be determined by computing the mean distance evolution for each subframe.

The video data may be received from a multi-modal camera including an infrared projector and sensor.

25       The subframes may comprise a regular grid of tiles (such as square tiles); however, many other arrangements are possible. The subframes may have many different shapes and it is not essential to all embodiments for all subframes to have the same shape.

30       The scene may be a fixed scene. For example, the scene may be captured using a camera having a fixed position and having unchanging direction and zoom.

The means may comprise: at least one processor; and at least one memory including computer program code, the at least one memory and the computer program code  
35       configured, with the at least one processor, to cause the performance of the apparatus.

In a second aspect, this specification describes a method comprising: receiving video data for a scene; determining a movement measurement for at least some of a plurality of subframes

of the video data; weighting the movement measurements to generate a plurality of weighted movement measurements, wherein the weighting is dependent on the subframe; and generating a movement indication for the scene from a combination (such as a sum) of some or all of the weighted movement measurements. The movement measurements for said  
5 subframes may comprise an indication of a degree of movement between successive instances of the subframes.

The weighting may be a binary weighting or a non-binary weighting. Alternatively, or in addition, the weighting for at least some of said subframes may be at least partially  
10 dependent on one or more of: historical movement data for that subframe; a location of the subframe within the scene; and visual characteristics of the scene.

In some embodiments, one or more of the weightings are updated on-the-fly/in use.

15 In some embodiments, determining the movement measurement for said subframes of the video data may comprise: determining a mean distance measurement for the respective subframe of the received video data; and determining, for one or more of the mean distance measurements, whether said mean distance measurement differs from a mean distance for a corresponding subframe of a preceding (e.g. immediately preceding) video data frame  
20 instance.

The movement measurements for said subframes of the video data may be determined on the basis of a change in received depth data for the respective subframe. The received depth data may, for example, be determined by computing the mean distance evolution for each  
25 subframe.

In a third aspect, this specification describes any apparatus configured to perform any method as described with reference to the second aspect.

30 In a fourth aspect, this specification describes computer-readable instructions which, when executed by computing apparatus, cause the computing apparatus to perform any method as described with reference to the second aspect.

In a fifth aspect, this specification describes a computer program comprising instructions for  
35 causing an apparatus to perform at least the following: receive video data for a scene; determine a movement measurement for at least some of a plurality of subframes of the video data; weight the movement measurements to generate a plurality of weighted movement measurements, wherein the weighting is dependent on the subframe; and generate a

movement indication for the scene from a combination (such as a sum) of some or all of the weighted movement measurements.

In a sixth aspect, this specification describes a computer-readable medium (such as a non-transitory computer readable medium) comprising program instructions stored thereon for performing at least the following: receiving video data for a scene; determining a movement measurement for at least some of a plurality of subframes of the video data; weighting the movement measurements to generate a plurality of weighted movement measurements, wherein the weighting is dependent on the subframe; and generating a movement indication for the scene from a combination (such as a sum) of some or all of the weighted movement measurements.

In a seventh aspect, this specification describes an apparatus comprising: at least one processor; and at least one memory including computer program code which, when executed by the at least one processor, causes the apparatus to: receive video data for a scene; determine a movement measurement for at least some of a plurality of subframes of the video data; weight the movement measurements to generate a plurality of weighted movement measurements, wherein the weighting is dependent on the subframe; and generate a movement indication for the scene from a combination (such as a sum) of some or all of the weighted movement measurements.

In an eighth aspect, this specification describes an apparatus comprising: a first input for receiving video data for a scene and/or an imaging device for obtaining video data for a scene; a first control module for determining a movement measurement for at least some of a plurality of subframes of the video data; a second control module for weighting the movement measurements to generate a plurality of weighted movement measurements, wherein the weighting is dependent on the subframe; and a third control module for generating a movement indication for the scene from a combination (such as a sum) of some or all of the weighted movement measurements. At least some of the first, second and third control modules may be combined.

### **Brief description of the drawings**

Examples will now be described, by way of non-limiting examples, with reference to the following schematic drawings, in which:

FIG. 1 is block diagram of a system in accordance with an example embodiment;  
FIG. 2 is a flow chart showing an algorithm in accordance with an example embodiment;  
FIG. 3 shows an example output in accordance with an example embodiment;

FIG. 4 shows an example output in accordance with an example embodiment;

FIG. 5 is block diagram of a system in accordance with an example embodiment;

FIG. 6 is a flow chart showing an algorithm in accordance with an example embodiment;

FIG. 7 is a flow chart showing an algorithm in accordance with an example embodiment;

5 FIG. 8 shows initialisation data in accordance with an example embodiment;

FIGS. 9A and 9B show results in accordance with an example embodiment;

FIGS. 10A and 10B show results in accordance with an example embodiment;

FIG. 11 is a block diagram of components of a system in accordance with an example embodiment; and

10 FIGS. 12A and 12B show tangible media, respectively a removable memory unit and a compact disc (CD) storing computer-readable code which when run by a computer perform operations according to example embodiments.

### Detailed description

15 Sleep disorders are associated with many problems, including psychiatric and medical problems. Sleep disorders are sometimes divided into a number of subcategories, such as intrinsic sleep disorders, extrinsic sleep disorders and circadian rhythm based sleep disorders.

20 Examples of intrinsic sleep disorders include idiopathic hypersomnia, narcolepsy, periodic limb movement disorder, restless legs syndrome, sleep apnea and sleep state misperception.

Examples of extrinsic sleep disorders include alcohol-dependent sleep disorder, food allergy insomnia and inadequate sleep routine.

25

Examples of circadian rhythm sleep disorders include advanced sleep phase syndrome, delayed sleep phase syndrome, jetlag and shift worker sleep disorder.

30 Better understanding of sleep physiology and pathophysiology may aid in improving the care received by individuals suffering with such difficulties. Many sleep disorders are primarily diagnosed based on self-reported complaints. Lack of objective data can hinder case understanding and care provided.

35 Non-contact, discrete longitudinal home monitoring may represent the most suitable option for monitoring sleep quality of a period of time and may, for example, be preferred to clinic-based monitoring as the sleep patterns can vary between days depending on food intake, lifestyle and health status.

- 6 -

FIG. 1 is block diagram of a system, indicated generally by the reference numeral 10, in accordance with an example embodiment. The system 10 includes a bed 12 and a video camera 14. The video camera may, for example, be mounted on a tripod 15. As discussed further below, the camera 14 may be a depth camera, although this is not essential to all  
5 embodiments.

In the use of the system 10, a patient (or some other subject) may sleep on the bed 12, with the video camera 14 being used to record aspects of sleep quality. The bed 12 may, for example, be in a patient's home, thereby enabling home monitoring, potentially over an  
10 extended period of time.

FIG. 2 is a flow chart showing an algorithm, indicated generally by the reference numeral 20, in accordance with an example embodiment.

15 The algorithm 20 starts at operation 22, where video data is received (for example from the camera 14 described above). In some example embodiments, the camera 14 is a depth camera (e.g. a camera incorporating an RGB-D sensor) such that the video data received in the operation 22 comprises depth data for a plurality of pixels of an image of a scene. An example  
20 depth camera may capture per-pixel depth information via infra-red (IR) mediated structured light patterns with stereo sensing or time-of-flight sensing to generate a depth map. Some depth cameras include three-dimensional sensing capabilities allowing synchronized image streaming of depth information. Many other sensors are possible (including video sensors that do not provide depth information).

25 FIG. 3 shows an example output, indicated generally by the reference numeral 30, in accordance with an example embodiment. The output 30 is an example of video data received in the operation 22 of the algorithm 20.

30 The video data received in operation 22 may be video data of a scene that is fixed, with the camera in a fixed position and having unchanging direction and zoom. Thus, for example, when a person is sleeping on the bed 12 of the system 10 described above, changes in the video data relate to movement of that person. By providing a fixed scene, changes in the video data (due, for example, to user movement) can be monitored over time (e.g. over several hours of sleep).

35 Movement measurements are determined at operation 24 of the algorithm 20. For example, movement between successive frames of the video data captured in operation 22 may be determined.

More specifically, the operation 24 may determine movement measurements for each of a plurality of subframes of the video data. By way of example, FIG. 4 shows an example output, indicated generally by the reference numeral 40, in accordance with an example

5 embodiment. The output 40 shows the data of the output 30 that has been divided into a plurality of subframes (sixteen subframes are shown in the output 40). A subframe is a sub-division of a frame of video data. For example, a frame may be divided into a plurality of subframes. The subframes may be square, but may take on a different form (such as circular shapes or irregular shapes). Moreover, it is not essential for all subframes to have the same  
10 shape.

Thus, although the example output 40, comprises subframes in the form of a regular grid of tiles, this is not essential. Other arrangements are possible. The size, shape and/or arrangement of subframes may be different and may be configurable.

15 The operation 24 may determine movement measurements between successive instances of each subframe. In one example implementation, the successive instances of the subframes were separated by 5 seconds (although, of course, other separations are possible). As discussed further below, the movement measurements for each of the plurality of subframes  
20 of the video data may be determined on the basis of a change in the received depth data for the respective subframe (e.g. by computing the mean distance evolution for each subframe). Separating successive instances of the subframes by a period such as 5 seconds even if more subframes are available may increase the likelihood of measurable movement occurring between frames and may reduce the data storage and processing requirements of the system.

25 The determined movement measurements are weighted at operation 26 to generate a plurality of weighted movement measurements. As discussed in detail below, the weighting may be dependent on the subframe. For example, the weighting for each subframe may be at least partially based on historical movement data for that subframe. Thus, subframes in  
30 which movement has occurred in the recent past may have a higher weighting than subframes in which such movement has not occurred.

Finally, a movement indication, based on the weighted movement measurement, is provided in operation 28. By way of example, the operation 28 may determine whether movement is  
35 detected in any of the subframes of an instance of the video image. The movement indication may be generated from a combination (such as a sum) of some or all of the weighted movement measurements.

FIG. 5 is block diagram of a system, indicated generally by the reference numeral 50, in accordance with an example embodiment. The system 50 comprises an imaging device 52 (such as the camera 14), a data storage module 54, a data processing module 56 and a control module 58.

5 The data storage module 54 may store measurement data (e.g. the output of the operation 24 of the algorithm 20) under the control of the control module 58, with the stored data being processed by data processor 56 to generate an output (e.g. the output of the operation 28 of the operation 20). The output may be provided to the control module 58. In an alternative  
10 arrangement, weighted data (e.g. the outputs of the operation 26) may be stored by the data storage module 54, with the stored (weighted) data being processed by data processor 56 to generate an output (e.g. the output of the operation 28 of the operation 20). A potential advantage of storing weighted data is that, in some circumstances, the storage requirements of the weighted data may be less than the storage requirements for the measurement data  
15 (particularly if, for example, the weighting is binary weighting).

As discussed further below, the video data received in the operation 22 is typically noisy. The weighting in operation 26 may seek to improve the signal-to-noise ratio of the data in order to seek to improve the quality of the movement indication generated in operation 28. A  
20 variety of weighting arrangements are described herein. Those weighting arrangements may be used alone or in any combination.

The weighting in operation 26 may be a binary weighting. Thus, some subframes may be ignored when generating the movement indication (in operation 26). With appropriate  
25 weighting, data from some subframes that do not contain useful movement data can be ignored, thereby increasing the overall signal-to-noise ratio. Data having a zero weighting may not need to be stored, potentially reducing both data storage and data processing requirements.

30 The weighting for a particular subframe may be at least partially dependent on historical movement data for that subframe. For example, subframes where movement has occurred in the past (e.g. in the recent past) may have a higher weighting. In this way, the weighted sum can be biased towards subframes in which movement has occurred in the past (e.g. the recent past). This tends to increase signal-to-noise ratio.

35 The weighting for a particular subframe may be at least partially dependent on a location of that subframe within the scene. For example, central subframes may have a higher weighting than peripheral subframes (on the assumption that movement is more likely to occur closer

to the centre of the image received in the operation 22). This may make sense since an operator of a camera (such as the camera 14) is likely to arrange the system such that a patient or subject on the bed 12 is close to the centre of the images captured by the camera.

- 5 The weighting for a particular subframe may be at least partially dependent on visual characteristics of the scene. For example, subframes with surfaces above the level of the bed 12 in the system 10 described above may be weighted higher than subframes with surfaces below the surface on the bed (on the assumption that movement is more likely to occur above the level of the bed than below the level of the bed).

10

The weighting for a particular subframe may change over time. For example, as more movement data for a particular subframe is obtained, the “historical movement data” that might be relevant to the weighting of that subframe may be updated on-the-fly.

- 15 The camera 14 of the system 10 and/or the imaging device 52 of the system 50 may be a multimodal camera comprising a colour (RGB) camera and an infrared (IR) projector and sensor. The sensor may send an array of near infra-red (IR) light into the field-of-view of the imaging device 52, with a detector receiving reflected IR and an image sensor (e.g. a CMOS image sensor) running computational algorithms to construct a real-time, three-dimensional  
20 depth value mesh-based video. Information obtained and processed in this way may be used, for example, to identify individuals, their movements, gestures and body properties and/or may be used to measure size, volume and/or to classify objects. Images from depth cameras can also be used for obstacle detection by locating a floor and walls.

- 25 In one example implementation, the imaging device was implemented using a Kinect (RTM) camera provided by the Microsoft Corporation. In one example, the frame rate of the imaging device 52 was 33 frames per second.

- 30 The IR projector and detector of the imaging device 52 may produce depth matrices, where each element of the matrix corresponds to the distance from an object in the image to the imaging device. The depth matrices can be converted to grey images (e.g. the darker the pixel, the closer it is to the sensor). If the object is too close or too far away, the pixels may be set to zero such that they appear black.

- 35 Depth frames may be captured by the imaging device 52 and stored in the data storage module 54 (e.g. in the form of binary files). In an example implementation, the data storage 54 was used to record depth frames for subjects sleeping for period ranging between 5 and 8

hours. In one example implementation, the data storage requirement for one night of sleep was of the order of 200GB.

The stored data can be processed by the data processor 56. The data processing may be online (e.g. during data collection), offline (e.g. after data collection) or a combination of the two. The data processor 56 may provide an output, as discussed further below.

FIG. 6 is a flow chart showing an algorithm, indicated generally by the reference numeral 60, showing an example implementation of the operation 24 of the algorithm 20, in accordance with an example embodiment. As described above, the operation 24 determines movement between successive frames of the captured video data.

The algorithm 60 starts at operation 62 where a mean distance value for each subframe is determined. Then, at operation 64, distances between subframes (based on the distances determined in operation 62) is determined. As discussed further below, a distance between subframes that is higher than a threshold value may be indicative of movement between successive frames.

By way of example, consider the video data 30 received in an example of the operation 22. A mean distance value between two frames of the video data 30 may be determined by comparing the mean distance value between two consecutive frames. In one example, all the pixels in the matrix were reordered in a vector of length 217088 (424\*512), with each row of the depth matrix being placed consecutively in the vector.

After removing all the zero pixels in the frames the mean distance of the frame number  $i$  is defined by:

$$M(i) = \frac{1}{N} \times \sum_{k=1}^N D(k, i) \quad (1)$$

Where:

$D \in G \setminus Z$  with

$G = \{\text{set of pixels in the frame}\};$  and

$Z = \{\text{set of zero value pixels in the frame}\}.$

Having computed the mean distance of successive frames of a dataset (the operation 62), a determination of movement in the relevant scene can be made (the operation 64).

One method for determining movement is as follows. First, a mean is defined based on the set of all the non-zero value pixels. Next, if the difference between the mean distance of two

consecutive frames is above a certain threshold  $\theta_1$ , a change in position is noticed. This is given by:  $M(i) - M(i - 1) \geq \theta_1$ .

Several values using the standard deviation and maximum difference are tested in order to find the best threshold values (as discussed further below).

Consider now the video data 40 received in an example of the operation 22. A mean distance value may be determined by determining the mean distance of each subframe.

The mean distance of sub frame  $j$  in frame number  $i$  is defined as:

$$M(i, j) = \frac{1}{N'} \times \sum_{k=1}^{N'} D(k, i, j) \quad (2)$$

Where:

$D \in G' \setminus Z'$  with

$G' = \{\text{set of pixels in the subframe } j\}$ ; and

$Z' = \{\text{set of zero value pixels in the subframe } j\}$ .

Now, instead of detecting the movements in the entire frame, we may detect movement in each subframe. If the mean distance between two consecutive subframes differs from a certain threshold  $\theta_2$  the algorithm will detect a movement in this sub frame and by consequence in the entire frame. This is given by:  $M(i, j) - M(i - 1, j) \geq \theta_2$ .

FIG. 7 is a flow chart showing an algorithm, indicated generally by the reference numeral 70, in accordance with an example embodiment. The algorithm 70 starts at operation 72 where an initialisation phase is carried out. The initialisation phase is conducted without a patient (or some other subject) being present (such that there should be no movement detected). As described below, the initialisation phase is used to determine a noise level in the data. Next, at operation 74, a data collection phase is carried out. Finally, at operation 76, subframes (of the data collected in operation 74) having data indicative of movement are identified (thereby implementing the operation 28 of the algorithm 20 described above). The subframes are identified in operation 76 based on a threshold distance level that is set depending on the noise level determined in the operation 72.

FIG. 8 shows an initialisation data, indicated generally by the reference numeral 80 in accordance with an example embodiment. The initialisation data 80 includes mean distance data 82, a first image 83, a second image 84 and a representation of the noise 85. The first image 83 corresponds to the mean distance data of a data point 86 and the second image 84 corresponds to the mean distance data of a data point 87.

The mean distance data 82 shows how the determined mean frame distance changes over time (a time period of 18 minutes is shown in FIG. 8). As the data 80 was collected without motion (e.g. without a patient or subject being on the bed 12 of the system 10), the variations in the mean distance data 82 are representative of noise. The noise representation 85 expresses the noise as a Gaussian distribution. That distribution can be used to determine a standard deviation for the noise data (as indicated by the standard deviation shown in the data 82). A maximum difference between two consecutive mean values is also plotted on the output 80.

The operation 72 seeks to assess the intrinsic noise in the frames of the output 80. In this way, the intrinsic variation in the data at rest can be determined so that the detection of motion will not be confused by noise. The operation 72 may be implemented using the principles of the algorithm 20 described above. Thus, video data (i.e. the output 82) may be received. Distances in the data 82 can then be determined (see operation 24 of the algorithm) and may be weighted (operation 26 of the algorithm 20).

In the data collection phase 74 of the algorithm 70, the subject goes to sleep and the relevant dataset is collected. Then, in operation 76, the knowledge of the noise is exploited to determine the likelihood that a change in distance data is indicative of actual movement (rather than noise). By way of example, movement may be deemed to have occurred in the event that distance data received in a subframe is more than five standard deviations (as determined in operation 72) away from the previous mean distance.

## Experimental Results

An experiment has been conducted on fourteen nights including eleven different subjects in order to validate the reproducibility of the experiment. The data were stored as videos which serve as ground truth. The true detection of movements was manually noted looking at the videos and was compared to the detection of movements provided by the algorithm being tested.

FIGS. 9A and 9B show results, indicated generally by the reference numerals 92 and 94 respectively, in accordance with example embodiments. The results are based on a mean distance evolution (i.e. how the mean distance determinations change over time) for the fourteen subjects. The percent of good detection (true positives) averaging all experiments is shown together with the overestimation of movements (false positives). The results 92 and 94 are based on frames of image data having the form described above with reference to FIG. 3 (i.e. without dividing the frames into subframes).

The results 92 show the number of correct movement identifications (true positives) and the number of false movement identification (false positives) for different standard deviation threshold levels. With the movement threshold at two standard deviations (such that a mean distance change of at least two standard deviations is detected), a true positives measurement of 45% was detected and a false positives measurements of 81% was detected. As the threshold level was increased, the number of false positives reduced (to zero at five standard deviations). However, at five standard deviations, the true positives level had reduced to 26%. Accordingly, the results 92 show poor performance.

The results 102 show the number of correct movement identifications (true positives) and the number of false movement identification (false positives) when using maximum distance differences as the threshold value. The performance was also poor.

FIGS. 10A and 10B show results, indicated generally by the reference numerals 102 and 104 respectively, in accordance with example embodiments. The results are based on a mean distance evolution for the fourteen subjects. The percent of good detection (true positives) averaging all experiments is shown together with the overestimation of movements (false positives). The results 102 and 104 are based a frames of image data having the form described above with reference to FIG. 4 (i.e. dividing the frames into subframes and considering each subframe separately).

The results 102 show the number of correct movement identifications (true positives) and the number of false movement identification (false positives) for different standard deviation threshold levels. With the movement threshold at five standard deviations, a true positives measurement of 70% was detected and a false positives measurement of 3% was detected. Thus, the results for the subframes arrangement were significantly better than the results for the full frame arrangement described above with reference to FIGS. 9A and 9B. The results 104 show a similarly good performance when using maximum distance differences as the threshold value.

The examples described above generally relate to the use of images from a depth camera in the generation of a movement indication for a scene. This is not essential to all embodiments. For example, changes in subframes of successive video images can be used to determine movement in a scene.

The examples described above generally relate to sleep monitoring in general, and the detection of data relating to sleep disorders in particular. This is not essential to all

embodiments. For example, the principles discussed herein can be used to determine movement in a scene for other purposes, for example for tracking movement in video streams for the purposes of detecting the presence of moving animals (i.e. “trail cameras” used in hunting or conservation) or for detecting a moving person for security monitoring purposes.

5

For completeness, FIG. 11 is a schematic diagram of components of one or more of the example embodiments described previously, which hereafter are referred to generically as processing systems 300. A processing system 300 may have a processor 302, a memory 304 closely coupled to the processor and comprised of a RAM 314 and ROM 312, and, optionally, user input 310 and a display 318. The processing system 300 may comprise one or more network/apparatus interfaces 308 for connection to a network/apparatus, e.g. a modem which may be wired or wireless. Interface 308 may also operate as a connection to other apparatus such as device/apparatus which is not network side apparatus. Thus direct connection between devices/apparatus without network participation is possible.

15

The processor 302 is connected to each of the other components in order to control operation thereof.

20

The memory 304 may comprise a non-volatile memory, such as a hard disk drive (HDD) or a solid state drive (SSD). The ROM 312 of the memory 314 stores, amongst other things, an operating system 315 and may store software applications 316. The RAM 314 of the memory 304 is used by the processor 302 for the temporary storage of data. The operating system 315 may contain code which, when executed by the processor implements aspects of the algorithms 20, 60 and 70 described above. Note that in the case of small device/apparatus the memory can be most suitable for small size usage i.e. not always hard disk drive (HDD) or solid state drive (SSD) is used.

25

The processor 302 may take any suitable form. For instance, it may be a microcontroller, a plurality of microcontrollers, a processor, or a plurality of processors.

30

The processing system 300 may be a standalone computer, a server, a console, or a network thereof. The processing system 300 and needed structural parts may be all inside device/apparatus such as IoT device/apparatus i.e. embedded to very small size

35

In some example embodiments, the processing system 300 may also be associated with external software applications. These may be applications stored on a remote server device/apparatus and may run partly or exclusively on the remote server device/apparatus. These applications may be termed cloud-hosted applications. The processing system 300

may be in communication with the remote server device/apparatus in order to utilize the software application stored there.

FIGS. 12A and 12B show tangible media, respectively a removable memory unit 365 and a compact disc (CD) 368, storing computer-readable code which when run by a computer may perform methods according to example embodiments described above. The removable memory unit 365 may be a memory stick, e.g. a USB memory stick, having internal memory 366 storing the computer-readable code. The memory 366 may be accessed by a computer system via a connector 367. The CD 368 may be a CD-ROM or a DVD or similar. Other forms of tangible storage media may be used. Tangible media can be any device/apparatus capable of storing data/information which data/information can be exchanged between devices/apparatus/network.

Embodiments of the present invention may be implemented in software, hardware, application logic or a combination of software, hardware and application logic. The software, application logic and/or hardware may reside on memory, or any computer media. In an example embodiment, the application logic, software or an instruction set is maintained on any one of various conventional computer-readable media. In the context of this document, a “memory” or “computer-readable medium” may be any non-transitory media or means that can contain, store, communicate, propagate or transport the instructions for use by or in connection with an instruction execution system, apparatus, or device, such as a computer.

Reference to, where relevant, “computer-readable storage medium”, “computer program product”, “tangibly embodied computer program” etc., or a “processor” or “processing circuitry” etc. should be understood to encompass not only computers having differing architectures such as single/multi-processor architectures and sequencers/parallel architectures, but also specialised circuits such as field programmable gate arrays FPGA, application specific circuits ASIC, signal processing devices/apparatus and other devices/apparatus. References to computer program, instructions, code etc. should be understood to express software for a programmable processor firmware such as the programmable content of a hardware device/apparatus as instructions for a processor or configured or configuration settings for a fixed function device/apparatus, gate array, programmable logic device/apparatus, etc.

As used in this application, the term “circuitry” refers to all of the following: (a) hardware-only circuit implementations (such as implementations in only analogue and/or digital circuitry) and (b) to combinations of circuits and software (and/or firmware), such as (as applicable): (i) to a combination of processor(s) or (ii) to portions of processor(s)/software

- 16 -

(including digital signal processor(s)), software, and memory(ies) that work together to cause an apparatus, such as a server, to perform various functions) and (c) to circuits, such as a microprocessor(s) or a portion of a microprocessor(s), that require software or firmware for operation, even if the software or firmware is not physically present.

5

If desired, the different functions discussed herein may be performed in a different order and/or concurrently with each other. Furthermore, if desired, one or more of the above-described functions may be optional or may be combined. Similarly, it will also be appreciated that the flow diagrams of Figure 2, 6 and 7 are examples only and that various operations depicted therein may be omitted, reordered and/or combined.

10

It will be appreciated that the above described example embodiments are purely illustrative and are not limiting on the scope of the invention. Other variations and modifications will be apparent to persons skilled in the art upon reading the present specification.

15

Moreover, the disclosure of the present application should be understood to include any novel features or any novel combination of features either explicitly or implicitly disclosed herein or any generalization thereof and during the prosecution of the present application or of any application derived therefrom, new claims may be formulated to cover any such features and/or combination of such features.

20

**Claims**

1. An apparatus comprising:  
means for receiving video data for a scene;  
5 means for determining a movement measurement for at least some of a plurality of subframes of the video data;  
means for weighting the movement measurements to generate a plurality of weighted movement measurements, wherein the weighting is dependent on the subframe; and  
means for generating a movement indication for the scene from a combination of  
10 some or all of the weighted movement measurements.
2. An apparatus as claimed in claim 1, wherein the weighting is a binary weighting.
3. An apparatus as claimed in claim 1 or claim 2, wherein the weighting for at least some  
15 of said subframes is at least partially dependent on one or more of:  
historical movement data for that subframe;  
a location of the subframe within the scene; and  
visual characteristics of the scene.
- 20 4. An apparatus as claimed in any one of the preceding claims, wherein one or more of the weightings are updated on-the-fly.
5. An apparatus as claimed in any one of the preceding claims, wherein the movement measurements for said subframes comprise an indication of a degree of movement between  
25 successive instances of the respective subframes.
6. An apparatus as claimed in any one of the preceding claims, wherein determining the movement measurement for said subframes of the video data comprises:  
means for determining a mean distance measurement for the respective subframes of  
30 the received video data; and  
means for determining, for one or more of the mean distance measurements, whether said mean distance measurement differs from a mean distance for a corresponding subframe of a preceding video data frame instance.
- 35 7. An apparatus as claimed in any one of the preceding claims, wherein the means for determining the movement measurements for said subframes of the video data determines that a movement has taken place if a mean distance measurement between successive subframes is greater than a threshold value.

8. An apparatus as claimed in any one of the preceding claims, wherein the video data comprises depth data for a plurality of pixels of an image of the scene.

5 9. An apparatus as claimed in claim 8, wherein the movement measurements for said subframes of the video data are determined on the basis of a change in received depth data for the respective subframe.

10 10. An apparatus as claimed in any one of the preceding claims, wherein the video data is received from a multi-modal camera including an infrared projector and sensor.

11. An apparatus as claimed in any one of the preceding claims, wherein the subframes comprise a regular grid of tiles.

15 12. An apparatus as claimed in any one of the preceding claims, wherein the scene is a fixed scene.

13. An apparatus as claimed in any one of the preceding claims, wherein the means comprise:

20 at least one processor; and

at least one memory including computer program code, the at least one memory and the computer program code configured, with the at least one processor, to cause the performance of the apparatus.

25 14. A method comprising:

receiving video data for a scene;

determining a movement measurement for at least some of a plurality of subframes of the video data;

30 weighting the movement measurements to generate a plurality of weighted movement measurements, wherein the weighting is dependent on the subframe; and

generating a movement indication for the scene from a combination of some or all of the weighted movement measurements.

35 15. A computer program comprising instructions for causing an apparatus to perform at least the following:



receive video data for a scene;


determine a movement measurement for at least some of a plurality of subframes of the video data;

- 19 -



weight the movement measurements to generate a plurality of weighted movement measurements, wherein the weighting is dependent on the subframe; and

generate a movement indication for the scene from a combination of some or all of the weighted movement measurements.









*FIG. 3*



*FIG. 4*



**FIG. 7**






FIG. 8

*FIG. 9A**FIG. 9B*




FIG. 10A







FIG. 10B

**FIG. 11**



*FIG. 12A*



*FIG. 12B*

# INTERNATIONAL SEARCH REPORT

International application No  
PCT/EP2019/083428

A. CLASSIFICATION OF SUBJECT MATTER  
INV. G06T7/00  
ADD.

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)  
G06T A61B

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

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

EPO-Internal, INSPEC, WPI Data

## C. DOCUMENTS CONSIDERED TO BE RELEVANT

| Category* | Citation of document, with indication, where appropriate, of the relevant passages                                                                                                                                                                                                                                    | Relevant to claim No. |
|-----------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------|
| X         | US 2018/064369 A1 (GUNTHER JACOB [US] ET AL) 8 March 2018 (2018-03-08)<br>abstract<br>figures 1A,2A<br>paragraphs [0040], [0041], [0053] - [0068]<br>-----                                                                                                                                                            | 1-15                  |
| X         | Nijholt P.J.: "Contactless respiration monitoring using a 3D camera system", Master thesis, 1 February 2016 (2016-02-01), pages 1-81, XP055562910,<br>Retrieved from the Internet:<br>URL:https://research.tue.nl/files/51211959/Nijholt_2016.pdf<br>[retrieved on 2019-02-28]<br>sections 4.6, 4.6.1, 5.2.3<br>----- | 1,14,15               |
| A         | -----<br>-/--                                                                                                                                                                                                                                                                                                         | 7                     |



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

7 January 2020

Date of mailing of the international search report

21/01/2020

Name and mailing address of the ISA/

European Patent Office, P.B. 5818 Patentlaan 2  
NL - 2280 HV Rijswijk  
Tel. (+31-70) 340-2040,  
Fax: (+31-70) 340-3016

Authorized officer

Scholz, Volker

# INTERNATIONAL SEARCH REPORT

International application No  
PCT/EP2019/083428

C(Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT

| Category* | Citation of document, with indication, where appropriate, of the relevant passages                                                                                                                                                                                                                                                                                                                                                    | Relevant to claim No. |
|-----------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------|
| A         | Jochen Kempfle: "Respiration Rate Estimation with Depth Cameras: An Evaluation of Parameters", Sensor-based Activity Recognition and Interaction, 20 September 2018 (2018-09-20), pages 1-10, XP055562919, 2 Penn Plaza, Suite 701New YorkNY10121-0701USA<br>ISBN: 978-1-4503-6487-4<br>Retrieved from the Internet:<br>URL:https://www.eti.uni-siegen.de/ubicomp/papers/ubi_iwoar2018c.pdf<br>[retrieved on 2019-02-28]<br>section 2 | 5,10                  |
| A         | -----<br>KAVEH BAKHTIYARI ET AL: "KinRes", PERVASIVE COMPUTING TECHNOLOGIES FOR HEALTHCARE, ACM, 2 PENN PLAZA, SUITE 701NEW YORKNY10121-0701USA, 23 May 2017 (2017-05-23), pages 472-475, XP058386283,<br>DOI: 10.1145/3154862.3154896<br>ISBN: 978-1-4503-6363-1<br>figure 4<br>section "Noise filtering"<br>-----                                                                                                                   | 7                     |

## INTERNATIONAL SEARCH REPORT

### Information on patent family members

International application No

PCT/EP2019/083428

| Patent document<br>cited in search report | Publication<br>date | Patent family<br>member(s) | Publication<br>date |
|-------------------------------------------|---------------------|----------------------------|---------------------|
| US 2018064369                             | A1                  | 08-03-2018                 | NONE                |
| -----                                     |                     |                            |                     |