



(51) International Patent Classification:
G06K 9/00 (2006.01)

(21) International Application Number:
PCT/GB2018/052478

(22) International Filing Date:
31 August 2018 (31.08.2018)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
1713977.5 31 August 2017 (31.08.2017) GB

(71) Applicant: CALIPSA LIMITED [GB/GB]; 8 Duncannon Street, London Greater London WC2N 4JF (GB).

(72) Inventors: GAO, Jiameng; Flat 46, Amisha Court, 161 Grange Road, London Greater London SE1 3GH (GB). PLOIX, Boris; 4 Linden Gardens, London Greater London W2 4ES (GB).

(74) Agent: BARNES, Philip; The IP Asset Partnership Limited, Prama House, 267 Banbury Road, Oxford Oxfordshire OX2 7HT (GB).

(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,

(54) Title: ANOMALY DETECTION FROM VIDEO DATA FROM SURVEILLANCE CAMERAS




Figure 2

(57) Abstract: The present invention relates to object orientated data analysis. More particularly, the present invention relates to analysis of objects within video data from surveillance cameras. According to a first aspect, there is provided a method of detecting anomalous behaviour, the method comprising the steps of: receiving a first set of input data, comprising one or more digital image frames; generating a statistical model based on the first set of input data, the statistical model operable to detect one or more objects within the first set of input data; analysing a second set of input data with respect to the statistical model; and detecting one or more objects within the second set of input data.



TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW,
KM, ML, MR, NE, SN, TD, TG).

Declarations under Rule 4.17:

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

Published:

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

ANOMALY DETECTION FROM VIDEO DATA FROM SURVEILLANCE CAMERAS

Field

5 The present invention relates to object orientated data analysis. More particularly, the present invention relates to analysis of objects within video data from surveillance cameras.

Background

10 It is estimated that there are close to 250 million video surveillance cameras deployed worldwide, capturing 1.6 trillion hours of video annually. In order to review even 20% of the most critical video streams, either in real time or post-processing, approximately 110 million human operators would be required to keep up. Human errors may also be present, and important details may be missed from the video stream.

15

Each human operator also requires training, and hence expansion of surveillance systems may be difficult to scale effectively. Further, humans are often not suited to watching video streams on multiple screens simultaneously for many hours. Focus will be lost, and important details become increasingly likely to be missed.

20

Summary of Invention

Aspects and/or embodiments seek to provide a method, apparatus and system to detect anomalous behaviour from video data.

25

According to a first aspect, there is provided a method of detecting anomalous behaviour, the method comprising the steps of: receiving a first set of input data, comprising one or more digital image frames; generating a statistical model based on the first set of input data, the statistical model operable to detect one or more objects within the first set of input data;
30 analysing a second set of input data with respect to the statistical model; and detecting one or more objects within the second set of input data.

Surveillance cameras are becoming increasingly popular. The videos they record can help catch criminals or detect certain behaviours. They may also be useful in detecting potential
35 accidents before they occur, and analysing behaviours and movements to ensure that such accidents do not happen in the future. However in order to do so, the video which is recorded needs to be viewed and analysed. By using the abovementioned method, such

analysis may be performed ceaselessly and with a lower rate of errors than a human camera operator. Human operators can be used to train and hence generate the statistical model as they would train a new employee. The model can therefore be arranged to learn relentlessly and is less likely to make the same mistake twice. By training over a wide range of video and
5 image data under different environmental conditions, algorithms used (which may comprise convolutional neural networks) can be arranged to be robust to a wide range of environment changes. A further advantage may be provided in that the same models may be deployed to various environments without further specific engineering and/or training. Any algorithms used may hence be robust to weather, lighting and camera placement issues and work
10 without any further configuration. This method may achieve levels of accuracy far beyond that of a human operator at a much lower cost. Typically, 50-80% may be saved using this method compared to manual enumeration methods.

Optionally, the first and/or second set of input data comprises one or more digital videos,
15 formed from the one or more digital image frames.

Digital videos are often difficult to analyse, as the task of watching them can be mentally unstimulating and hence not performed effectively. Humans are not generally suited to watching endless hours of video on multiple screens, as they get tired and lose focus.
20 However using the method disclosed herein, any human operators may instead be provided with actionable alerts as opposed to raw video feeds, keeping them more engaged and making the best use of their decision-making skills. A statistical model does not tire or require breaks, and can be able to process several video streams simultaneously.

25 Optionally, the one or more digital videos are recorded from one or more surveillance cameras.

Surveillance cameras, while often placed at sites of interest, often fail as the videos which they record are not fully analysed. Such a system would allow the video produced by the
30 surveillance cameras to be used to its full effect. The method disclosed herein is agnostic to the type of camera used to record video, for example regardless of whether the video was recorded with a mobile phone or a HD video recorder, the method disclosed herein may still be used in the video analysis.

35 Optionally, the generation of the statistical model is performed using one or more of: Convolutional Neural Networks (CNNs); Deep Convolutional Networks; Recurrent Neural

Networks; Reinforced Learning; Scale-Invariant Feature Transform (SIFT) features, and/or optical flow feature vectors.

5 CNNs, and/or similar tools, can provide a useful and robust means for generating a statistical model. They can be trained effectively, and learn from previous errors. This machine learning method may be used in combination with proprietary datasets to deliver greater accuracy of analysis. However other arrangements may be used, for example hand designed or modelled filters, such as Scale-Invariant Feature Transform (SIFT) features, and/or optical flow feature vectors. CNNs themselves may comprise a series of layers (also
10 referred to as “sets”) of filters, wherein each set of filters may be of the same width and height. Each layer of filters can be convolved over its input and the outputs fed into a subsequent layer. Optionally, between each layer, the outputs are fed into a non-linear function before being fed into a subsequent layer. The values of the filters may be randomly selected at the beginning of a training session, and during gradient descent training, they
15 converge to optimal values in relation to the tasks for which they are being trained.

Optionally, the method further comprises the steps of: analysing the first set of input data through one or more filters; and obtaining one or more filter outputs. Optionally, the generation of the statistical model comprises the use of the one or more filter outputs.
20 Optionally, the one or more filters comprise one or more of: CNNs; Deep Convolutional Networks; Recurrent Neural Networks; Reinforced Learning; Scale-Invariant Feature Transform (SIFT) features, and/or optical flow feature vectors.

Filters can reduce noise and other distraction from a video and create higher-level
25 abstractions of images and video, hence enable a more efficient and accurate analysis. CNNs and/or similar tools may be used in place of one or more filters.

Optionally, the one or more objects comprise one or more of: vehicles; human beings; animals; plants; buildings; and/or weather formations. Optionally, the statistical model is
30 operable to track one or more objects in the first and/or second set of input data. Optionally, the statistical model is operable detect anomalous objects in the first and/or second set of input data.

It can be advantageous to analyse the presence of certain objects. For example, accidents
35 can be detected through anomalous behaviours, for example collisions, between two vehicles. Unusual movements of human beings, for example a single person travelling against a large crowd, may also be useful to detect. Such analysis may allow accidents to

be prevented in the future, or an emergency team to be dispatched to the site of an accident more efficiently.

Optionally, the analysis of the second set of input data is unsupervised. Optionally, the
5 analysis of the second set of input data occurs in real time.

By analysing a set of data without supervision, the analysis of the data may be rapidly scaled up. If a statistical model is trained to recognise anomalous behaviour on roads with an accuracy above a specified level, then video feeds from a large number of other surveillance
10 cameras may be immediately analysed using that same statistical model. Such a task would have previously required the employment and training of a correspondingly large number of people to review and label the raw video data. In this embodiment, instead of building a classification model that classifies the behaviour of objects into pre-determined classes of activities (e.g. car driving slowly, making illegal turns), a probability distribution may be
15 assigned over the probability of the positions and other features of the object. Therefore an anomaly may be detected by a probability value using a predetermined threshold.

According to a further aspect, there is provided an apparatus for detecting anomalous behaviour, comprising: means for receiving a first set of input data, comprising one or more
20 digital image frames; means for generating a statistical model based on the first set of input data, the statistical model operable to detect one or more objects within the first set of input data; means for analysing a second set of input data with respect to the statistical model; and means for detecting one or more objects within the second set of input data. According to a further aspect, there is provided a system operable to perform the method disclosed
25 herein. According to a further aspect, there is provided a computer program product operable to perform the method and/or apparatus and/or system disclosed herein.

By providing such an apparatus and/or system, the method disclosed herein may be effectively implemented. Such an implementation may allow for more effective use of
30 human resources within a company or organisation, as well as reducing undesirable anomalous behaviours and detecting their root causes.

Brief Description of Drawings

Embodiments will now be described, by way of example only and with reference to the
35 accompanying drawings having like-reference numerals, in which:

Figure 1 shows an example analysis of a video frame; and

Figure 2 shows a diagrammatic process flow.

5 Specific Description

Referring to Figure 1, a first embodiment will now be described. In this figure, a view 100 from a surveillance camera is shown. The view 100 comprises a single digital image frame taken from a plurality of such frames, together comprising a digital video. The surveillance camera from which this view 100 is observed is positioned over a road 110, and positioned in such a way that objects on the road 110 are visible. Such objects may comprise vehicles, road markings, vehicle paths, and road barriers. In this example, vehicles are being detected and tracked.

15 In particular, vehicles are being analysed in reference to what type of vehicle they appear to be. Cars 105 on the road 110 are identified as being different from motorcycles 115. The paths of the vehicles 120 are also identified and recorded. Anomalous interactions between vehicles, for example a collision, may be detected and recorded. Further, illegal manoeuvres such as forbidden lane changes and travelling at excessive speeds may be accurately monitored and recorded.

An exemplary analysis process is shown in Figure 2. In this figure, a first surveillance camera 205 records a video. This video may include, for example, a road on which traffic travels. The video is processed into an appropriately sized video file 210. The video file 210 is then used to develop a statistical model 215. For example, objects which may fall under the umbrella term of “vehicle” within the video file 210 may be identified by an operator, such that the statistical model 215 develops the ability to detect vehicles autonomously. Further, the statistical model 215 may be arranged to differentiate between different types of vehicle, for example categorising vehicles as “motorcycles”, “bicycles”, “cars”, “vans”, and/or “lorries”. This training is not necessarily limited to road surveillance videos, but could be applied to any video file. For example, surveillance camera video footage from a sports event could be used and the statistical model trained to recognise “humans”, or even “home fans” and “away fans”. If the detected humans were acting in an anomalous manner, for example a fight begins, the statistical model could be trained to recognise such behaviour. A simple self-service web interface may be used to train the statistical model 215 for personal use, for example in the case of a household camera used to detect cats in a garden, or differentiating black taxis from private cars on the street.

Once trained in such a manner, the statistical model 215 may then be applied to a second video file 211, derived from a second surveillance camera 206. The statistical model 215 may then autonomously detect the objects which it has been trained to recognise, which in this embodiment is vehicles travelling along a road. Such detection can occur without human supervision, and may be applied to many video feeds simultaneously, both recorded and in real time. The accuracy of the statistical model 215 may improve over time with supervision from a human operator. The algorithms used may be rewarded for correct notifications and penalised for false alarms.

A specific example of such an algorithm could comprise calculating a function $Q(s, a)$, where the expected reward Q is determined by the state "s" of the current conditions of the image or object, such as its position, its length of stay in the video, appearance features, etc. and actions "a", which would be to raise an alarm, or not raise an alarm. In such a scenario, the algorithm could take action a^* which has the maximum expected reward for a given state s . With such an arrangement in place, every time an alert is raised, or a notable event is missed, the user can reward the system for correct detections or penalise the system for false alarms or false negatives, which will reinforce and/or correct the function Q . Such correction may be arranged through gradient descent. Here the function Q may take the form of various differentiable function approximators, such as a neural network, with the state s and action a as input, or an individual Gaussian Process for each action a .

An output may be provided in the form of an annotated display 220, on which the analysed video file 211 is provided along with an overlay. The overlay can represent the findings of the statistical model, for example providing labels to any vehicles labelled "bicycle", or highlighting areas in which collisions seem to occur most frequently. An alert system may be employed to alert a human operator if a dangerous or otherwise anomalous situation arises. A web interface may be provided, comprising a search and reporting functionality. The interface may be operable to allow users to filter by event type, location and action taken. Interactive reports may be provided, with give a level of detail that is significantly more difficult to achieve with manual enumeration. For example, such a report may comprise vehicle speeds, classification, colour, changes in lanes, heatmaps and/or advanced flow analysis. The trained statistical model 215 may be stored in cloud storage, at a location remote from the site of the second surveillance camera 206. Any camera feeds or recorded videos may be uploaded to such a cloud platform using a simple interface. A user would only need access to a web browser and a working internet connection. No dedicated hardware

would be required in this case. Local analysis may also be provided, for example if privacy of the data was a major concern.

The implementation of the abovementioned arrangement may involve the following steps:

5

1. Obtain a sequence of images from a live or stored video feed, where the length of the sequence is set to a threshold. For example, 1 minute of video may be provided at a frame rate of 25 frames per second.

10

2. Pass the sequence of images as input to a convolutional neural network-based object detector, which in this embodiment comprises a convolutional neural network (CNN). The object detector may further or in addition comprise the use of at least one of: Deep Convolutional Networks; Recurrent Neural Networks; Reinforced Learning; Scale-Invariant Feature Transform (SIFT) features, and/or optical flow feature vectors. The CNN comprises a set of automatically learnt filters and produces a tensor for each image, and a region proposal component. The regional proposal component may be in the form of a neural network (a set of automatically learnt matrix transformations) that uses the tensor as input to estimate bounding box coordinates for regions with high likelihoods of objects. This gives, for each image, a set of bounding boxes for each of the detected objects

15

20

3. Using the sequence of images and the sets of bounding boxes, link the detection of each individual object together using a CNN-based tracker. A CNN-based tracker may comprise a CNN (itself comprising a collection of filters), which takes a crop of the image as input, wherein the crop may be centred on the target tracked object. The filter outputs may then be arranged to produce a tensor-representation which can be flattened to form a vector. For the next frame, the CNN may be applied across a search area around a previous location of the object, obtaining one or more tensor and/or vector outputs. A normalised dot-product may then be used to calculate the cosine between the new vectors and the vector of the target object. The location of the target object may then be assigned to be the location that produced the vector with the highest cosine value). Therefore by using a CNN (which could be different from the CNN used in the object detector) to produce a tensor representation for each object on the first frame (also referred to as the target objects) and matching, on a subsequent frame, detection bounding boxes which produce tensor representations may be generated that are the closest to the target object.

25

30

35

4. These bounding boxes may be drawn onto the sequence of images and displayed to the user in video form. The tracker may be arranged to output four coordinates (x_{minimum} , y_{minimum} , x_{maximum} , y_{maximum}). From these coordinates four lines may be generated to form a rectangle. Once one or more rectangles have been formed, a tool can be used to draw them onto an image. In this embodiment, such a tool may comprise a program arranged to change one or more colours of the pixels on the bounding box to a desired colour for the bounding box. Hence the tool may change the pixel values of the lines between (x_{minimum} , y_{minimum}) to (x_{maximum} , y_{minimum}), (x_{minimum} , y_{minimum}) to (x_{minimum} , y_{maximum}), etc. to red and/or blue and/or green.
5. Having done this for all detected objects through the sequence of images, the full trajectories of each of the detected objects may be obtained. This provides the 2D spatial coordinates for each object at every frame it is detected within the video. A set of 3D coordinates is therefore provided for each object, with the third dimension being time.
6. Using these 3D trajectory coordinates, the starting frame-number of each of the objects is set to zero, such that the third dimension represents how long the object has remained within the video.
7. Build a probability density model over the sets of 3D coordinates for all detected objects. This can be done using Kernel Density Estimation, where the kernel could be Gaussian, triangular, or any other suitable arrangement, such that the density at any point in this 3D space would correlate with the number of individual trajectory data points close by. This provides a method to numerically calculate the probability at any point for future data points.
8. For further images or sequences of images from the live or stored video feed, the same object detection and tracking methods may be used as previously described to obtain the 3D coordinates for each new object in the new images.
9. These coordinates may then be evaluated using the Kernel Density Estimation model. Conventionally, a Kernel Density Estimation model calculates the probability of a certain point by counting and/or weighting all nearby points using a function ("kernel"), such as a Gaussian or a triangular function, giving a weighted average for the expected number of data points at a predetermined location. For any data points

with a probability less than a threshold (for example, 0.05), an alert would be raised to the user. This could be extended to the user rewards and/or reinforcement learning arrangement as disclosed above. Examples of alerts include drawing the bounding box of the object in a different colour to distinguish them to the user, displaying a text box directly to the user, or sending the user an email.

A further embodiment is also disclosed as in the following steps:

1. Obtain a sequence of images from a live or stored video feed, where the length of the sequence is set to a threshold, as previously described.
2. Pass the sequence of images through a CNN (which could be different from the CNNs used in the object detector and tracker), such that a tensor representation is obtained for each whole image, where these tensors may have dimensions such as (40 X 30 X 512).
3. The tensor representation is then flattened into a vector (which in this exemplary embodiment is of dimension $614,400 = 40 \times 30 \times 512$), wherein this vector representation is used as the input for a probability density estimation model, an example of which is Kernel Density Estimation.
4. For all future frames, their tensor representations are flattened and their probability values may be estimated using the density estimation model, such that low probability scores (such as less than 0.05) would raise an alert to the user.

A yet further embodiment may at least in part combine the above methods in the following steps:

1. Obtain a sequence of images from a live or stored video feed, where the length of the sequence is set to a threshold, as previously described.
2. Pass the sequence of images through a CNN-based object detector and an object tracker, such that for each object, a trajectory (which may be in the form of a sequence of bounding boxes) is obtained for each object.
3. Make a crop of the image from each bounding box of each object, thus obtaining a set of the sequence of crops for each object as it moves through the frame.

4. This set of crops is then fed into a CNN (which could be different from the CNNs used in the object detector and tracker), wherein a tensor representation is obtained for each crop. All the obtained tensor representations are then flattened and used as input for a probability density estimation model, as described previously.
5. For future frames, the object detector and tracker are run on each new incoming frame, such that bounding boxes for each object are obtained. Cropped images of these detected objects are then fed into the CNN of step 4, to obtain flattened vector representations which are then fed into the density estimation model, where probability scores are obtained.
6. An alert is then sent to the user, as described above.

Any system feature as described herein may also be provided as a method feature, and vice versa. As used herein, means plus function features may be expressed alternatively in terms of their corresponding structure.

Any feature in one aspect of the invention may be applied to other aspects of the invention, in any appropriate combination. In particular, method aspects may be applied to system aspects, and vice versa. Furthermore, any, some and/or all features in one aspect can be applied to any, some and/or all features in any other aspect, in any appropriate combination.

It should also be appreciated that particular combinations of the various features described and defined in any aspects of the invention can be implemented and/or supplied and/or used independently.

CLAIMS:

1. A method of detecting anomalous behaviour, the method comprising the steps of:
receiving a first set of input data, comprising one or more digital image
5 frames;
generating a statistical model based on the first set of input data, the
statistical model operable to detect one or more objects within the first set of input
data;
10 analysing a second set of input data with respect to the statistical model; and
detecting one or more objects within the second set of input data.
2. A method as claimed in claim 1, wherein the first and/or second set of input data
comprises one or more digital videos, formed from the one or more digital image
frames.
- 15 3. A method as claimed in any one of claims 1 or 2, wherein the one or more digital
videos are recorded from one or more surveillance cameras.
4. A method as claimed in any preceding claim, wherein the generation of the statistical
20 model is performed using one or more of: Convolutional Neural Networks (CNNs);
Deep Convolutional Networks; Recurrent Neural Networks; Reinforced Learning;
Scale-Invariant Feature Transform (SIFT) features, and/or optical flow feature
vectors.
- 25 5. A method as claimed in any preceding claim, further comprising the steps of:
analysing the first set of input data through one or more filters; and
obtaining one or more filter outputs.
6. A method as claimed in claim 5, wherein the generation of the statistical model
30 comprises the use of the one or more filter outputs.
7. A method as claimed in any one of claims 5 or 6, wherein the one or more filters
comprise one or more of: Convolutional Neural Networks (CNNs); Deep
Convolutional Networks; Recurrent Neural Networks; Reinforced Learning; Scale-
35 Invariant Feature Transform (SIFT) features, and/or optical flow feature vectors.

8. A method as claimed in any preceding claim, wherein the one or more objects comprise one or more of: vehicles; human beings; animals; plants; buildings; and/or weather formations.
- 5 9. A method as claimed in any preceding claim, wherein the statistical model is operable to track one or more objects in the first and/or second set of input data.
10. A method as claimed in any preceding claim, wherein the statistical model is operable detect anomalous objects in the first and/or second set of input data.
- 10 11. A method as claimed in any preceding claim, wherein the analysis of the second set of input data is unsupervised.
12. A method as claimed in any preceding claim, wherein the analysis of the second set of input data occurs in real time.
- 15 13. An apparatus for detecting anomalous behaviour, comprising:
means for receiving a first set of input data, comprising one or more digital image frames;
20 means for generating a statistical model based on the first set of input data, the statistical model operable to detect one or more objects within the first set of input data;
means for analysing a second set of input data with respect to the statistical model; and
25 means for detecting one or more objects within the second set of input data.
14. A system operable to perform the method of any one of claims 1 to 12.
15. A computer program product operable to perform the method and/or apparatus and/or system of any preceding claim.
- 30




Figure 1




Figure 2

INTERNATIONAL SEARCH REPORT

International application No
PCT/GB2018/052478

A. CLASSIFICATION OF SUBJECT MATTER
INV. G06K9/00
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)
G06K

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	WO 2015/001544 A2 (AGENT VIDEO INTELLIGENCE LTD [IL]) 8 January 2015 (2015-01-08) page 1, line 2 - line 20 page 2, line 3 - line 9 page 2, line 26 - page 3, line 26 page 4, line 5 - line 19 page 6, line 1 - line 16 page 7, line 22 - line 27 page 10, line 21 - line 30 page 11, line 1 - line 27 page 12, line 5 - line 12 page 13, line 19 - page 14, line 6 page 15, line 19 - line 30 page 17, line 5 - line 15 page 21, line 3 - line 14 page 22, line 19 - line 21 page 22, line 31 - page 24, line 2 figures 1,3,4,5 -/--	1-15

Further documents are listed in the continuation of Box C.

See patent family annex.

* Special categories of cited documents :

<p>"A" document defining the general state of the art which is not considered to be of particular relevance</p> <p>"E" earlier application or patent but published on or after the international filing date</p> <p>"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)</p> <p>"O" document referring to an oral disclosure, use, exhibition or other means</p> <p>"P" document published prior to the international filing date but later than the priority date claimed</p>	<p>"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</p> <p>"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</p> <p>"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</p> <p>"&" document member of the same patent family</p>
---	---

Date of the actual completion of the international search 16 November 2018	Date of mailing of the international search report 27/11/2018
---	--

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 Moreno, Marta
--	---

INTERNATIONAL SEARCH REPORT

International application No
PCT/GB2018/052478

C(Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT		
Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	<p style="text-align: center;">-----</p> <p>WO 2010/111748 A1 (UNIV CURTIN TECH [AU]; VENKATESH SVETHA [AU]; SAHA BUDHADITYA [AU]; LA) 7 October 2010 (2010-10-07) paragraph [0002] paragraph [0006] - paragraph [0009] paragraph [0017] paragraph [0024] - paragraph [0025] paragraph [0027] paragraph [0058] paragraph [0129] figure 15</p>	1-4,8, 10-14
X,P	<p style="text-align: center;">-----</p> <p>GB 2 554 948 A (CALIPSA LTD [GB]) 18 April 2018 (2018-04-18) page 1, line 5 - line 7 page 5, line 19 - page 7, line 5</p> <p style="text-align: center;">-----</p>	1-6,8-15

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No

PCT/GB2018/052478

Patent document cited in search report	Publication date	Patent family member(s)	Publication date	
WO 2015001544	A2	08-01-2015	EP 3017403 A2	11-05-2016
			US 2016139977 A1	19-05-2016
			WO 2015001544 A2	08-01-2015

WO 2010111748	A1	07-10-2010	AU 2010230857 A1	20-10-2011
			CN 102449660 A	09-05-2012
			EP 2415022 A1	08-02-2012
			JP 5625039 B2	12-11-2014
			JP 2012523025 A	27-09-2012
			SG 174519 A1	28-10-2011
			US 2012063641 A1	15-03-2012
			WO 2010111748 A1	07-10-2010

GB 2554948	A	18-04-2018	NONE	
