



- (51) International Patent Classification:
G06T 7/00 (2006.01)
- (21) International Application Number:
PCT/EP2016/063650
- (22) International Filing Date:
14 June 2016 (14.06.2016)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:
15174937.1 2 July 2015 (02.07.2015) EP
- (71) Applicant: CONTINENTAL AUTOMOTIVE GMBH
[DE/DE]; Vahrenwalder Straße 9, 30165 Hannover (DE).
- (72) Inventors: GUERREIRO, Rui; Coventry Street 43A,
Brighton Sussex BN1 5PP (GB). KADYROV, Alexander;
18 Willowfield Road, Eastbourne Sussex BN22 8AL (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, DK, DM,
DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT,
HN, HR, HU, ID, IL, IN, IR, IS, JP, KE, KG, KN, KP, KR,
KZ, LA, LC, LK, LR, LS, 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: DETECTION OF LENS CONTAMINATION USING EXPECTED EDGE TRAJECTORIES




Fig. 3

(57) Abstract: A method for detecting contaminations on a lens of a vehicle camera. Transition values are assigned to pixels of image frames. Three-dimensional coordinates are assigned to the pixels by using a simplified model of the scene. Using the assigned coordinates and movement data, trajectories are assigned to the pixels. An expectation value is derived from a transition value of a starting pixel of a trajectory and a satisfied transition value is derived from transition values along the trajectory. The expectation value and the satisfied transition value are accumulated over time, and a cleanliness value is derived from a ratio of the accumulated satisfied transition value to the accumulated expectation value.

DETECTION OF LENS CONTAMINATION USING EXPECTED EDGE TRAJECTORIES

Present day passenger vehicles are increasingly equipped with
5 camera-based assistance systems, which capture the surrounding environment of the vehicle, and provide a variety of functions for improving driving safety and comfort. The functionality of these systems is based on the analysis of the recorded image data. Therefore, the quality of the system's
10 predictions is directly related to the quality of the image data.

One factor that affects the image quality and which is difficult to control is the degree of contamination of the optical
15 system of the camera. The cameras may be positioned at places with less risk of contamination, or the cameras may be cleaned by an electric wiper. Despite of these provisions, it is impossible to avoid a contamination of the optical system completely. Therefore, it has been proposed to detect dirt
20 particles on a camera lens automatically in order to trigger an appropriate action. An example for such an automatic detection of lens contaminations is disclosed in the European patent application EP 2351351.

25 The present specification discloses a method and corresponding device for detecting contaminations on a lens of a vehicle camera.

Image data, which comprises image frames is received from the
30 vehicle camera or from multiple vehicle cameras over a data connection, such as an automotive data bus or a data cable of the camera.

Transition values are assigned to pixels of the image frames using a transition indicator. The transition values indicate whether there is a clearly recognizable intensity transition from dark to bright or from bright to dark at a corresponding pixel location of the frame. The transition may in turn indicate that there is a corner, an edge or another type of feature point at the corresponding location.

A transition indicator is an operator or a filter, such as a Laplace detector or a Harris corner detector, which is applied to the pixel values of the frame and generates transition values, which are associated to the pixels of the image. Herein, a pixel may also stand for a low resolution pixel, which corresponds to multiple adjacent pixels of the frame. The operator or filter can be provided by a function of the pixels in the vicinity of a pixel, which assigns a transition value to the pixel. In particular, the function can be a linear function of the pixel values, which can be represented by a matrix, such as a convolution matrix.

Three-dimensional coordinates are assigned to the pixels using a model of the exterior scene. For example, in a simple model the scene is divided into a ground surface and into regions comprising vertical objects with surfaces that are aligned in parallel to the road boundaries, such as houses and walls at the side of the road. Further information, such as surface recognition via vanishing lines or stereo images may be used to improve the accuracy of the three-dimensional coordinates.

The assigned three-dimensional coordinates and movement data of the vehicle is used to assign trajectories to the pixels, the movement data being provided by one or more vehicle sensors, such as CAN data comprising velocity and steering angle, camera data, ego-motion data which is derived from the camera data, etc. In particular, the ego-motion data may be derived from cameras of a surround view system. In a simple example, a trajectory can be defined by only two points, such as a starting point (x_0, y_0) and an end point (x, y) .

According to pre-determined criteria such as criteria for the accuracy of the estimated motion vector, it is decided whether a motion vector is assigned to a pixel. For those pixels and frames to which a trajectory is assigned, an expectation value and a satisfied transition value are computed.

The expectation value is derived from a transition value of a starting pixel of the trajectory. The trajectory may extend forwards or backwards in time and it comprises pixels of a set of adjacent image frames.

The satisfied transition value is derived from transition values along the trajectory, which begins with the starting pixel and extends over a set of frames that are adjacent to the frame of the starting pixel.

The expectation value and the satisfied transition value are accumulated over time, wherein time is represented by the frame index. A cleanliness value of a corresponding location of the camera lens is derived from a ratio of the accumulated

satisfied transition value to the accumulated expectation value.

The cleanliness value, which is derived from accumulated quantities, and the accumulated quantities themselves may have a well-defined value, even if the pixel location in the current frame does not have an associated trajectory. If the available information is not sufficient to update the accumulated quantities, the cleanliness value is computed based on the previous values, which were obtained as a result of the last update.

In particular, the accumulation over time may comprise update rules according to which a new accumulated value is derived from a weighted average of the previous accumulated value and a new value of the expectation value or of the satisfied transition value. The weighting factor is determined by the magnitude of the transition value at the starting point of the trajectory. A stronger transition value results in - or contributes to - a higher weight for the new value and a smaller weight for the previous value.

According to a further embodiment, the expectation value $Tot(x, y)$ and the satisfied transition value $Sat(x, y)$ of a pixel location (x, y) are updated based on at least two quantities, wherein the quantity is chosen which has the highest value, i.e. which gives the strongest indication.

The first quantity indicates if there is a clearly visible transition at the pixel location (x, y) and the second quantity indicates whether a transition is found at (x, y) that

corresponds to the transition at a starting point (x_0, y_0) of the trajectory.

According to one embodiment, the first value is computed as a weighted sum of the previous accumulated satisfied transition value and the value one, if the transition value at the pixel is equal to one, and it is set to zero, if the transition value is less than one. Setting the first value to zero is equivalent to disregarding the value.

According to a further embodiment, the expectation value $Tot(x, y)$ and the satisfied transition value $Sat(x, y)$ of a pixel location (x, y) are updated based on at least three quantities, wherein the quantity is chosen which has the highest value, i.e. which gives the strongest indication.

The first quantity indicates if there is a clearly visible transition at the pixel location (x, y) , the second quantity indicates whether a transition is found at (x, y) that corresponds to the transition at a starting point (x_0, y_0) of the trajectory with respect to a trajectory that progresses forwards in time, the third quantity indicates whether a transition is found at (x, y) that corresponds to the transition at a starting point (x_0, y_0) of the trajectory with respect to a trajectory that progresses backwards in time.

According to one embodiment, only cleanliness values are considered which are based on a sufficiently high expectation value. The accumulated expectation value $Tot(x, y)$ is compared with a predetermined threshold. In particular, the threshold may be 0.9 or higher.

If the expectation value is greater than the pre-determined threshold, the cleanliness value $C(x, y)$ of the pixel is the ratio of the accumulated satisfied transition value with the accumulated expectation value. If the expectation value is less than the pre-determined threshold, the cleanliness value $C(x, y)$ of the pixel is represented as "undetermined", for example by assigning no value or by assigning a value representing the status "undetermined".

In particular, the transition values can be computed as a two dimensional convolution of a filter with the pixels of an image frame and assigning the resulting transition value to the respective high resolution or low resolution pixels. Pixels at the image boundaries may be disregarded, a one-sided transition value can be assigned to the boundary pixels or a suitable outer boundary may be provided, such that the boundary pixels become inner pixels. Specifically, a filter can be chosen which provides a discrete approximation to a magnitude of a Laplace operator.

In order to identify the clearly visible transitions, the computation of the transition values may furthermore comprise comparing the pixel transition values with a detection threshold and setting the respective transition value to one if the respective transition magnitude value is greater than the detection threshold. Thereby, the transition value becomes a value between 0 and 1 and can be interpreted as a detection probability.

If a threshold is provided and a value exactly meets the threshold, the output value depends on the implementation. Often, a result is not sensitive to exact values but only to

ranges of values. In this case, the output value can be the same as for values above or below the threshold or it could be some value in between.

- 5 In order to speed up the computation and to relax the requirements on the trajectory computation, the method may furthermore comprise the computation of low-resolution transition values of low-resolution pixels. As a consequence, the further computation steps, such as the update equations, the
10 time accumulation and the computation of the cleanliness value are based on the low-resolution pixels.

The computation of the low-resolution transition values comprises dividing the image frames into disjoint subregions and
15 computing low resolution transition values by selecting maximum values of the transition values in the respective subregions. The low resolution transition values are stored in a memory, such as a circular buffer, for a predetermined number of image frames. The circular buffer is also referred to as
20 ring buffer of "FIFO buffer". The low resolution transition values are then used to compute the corresponding expectation values and the corresponding satisfied transition values.

The movement data may be provided in various ways. According
25 to one embodiment, the movement data is provided by a vehicle speed sensor and a steering angle sensor. According to another embodiment, the movement data is provided by an ego-motion computation, which is derived from at least one vehicle camera sensor.

30

The assignment of motion vectors, or of a trajectory, to the pixels is done using a simplified model of a surrounding

scene. According to one embodiment, the simplified model comprises a horizontal ground surface, a left region of vertically oriented objects and a right region of vertically oriented objects.

5

The prediction of a trajectory is particularly simple for a straight movement of a car. According to one embodiment sets of nearby frames are identified for which the vehicle is moving essentially in a straight line, or in a straight line within a predetermined margin. The accumulated expectation value and the accumulated satisfied transition value are then updated for the identified frames only.

10

In a further aspect, the present specification discloses a computer program product with computer readable code for executing the aforementioned method.

15

According to a further aspect, the present specification discloses an evaluation unit. The evaluation unit comprising an input connection for receiving image data and a circular buffer for storing a predetermined number of image frames.

20

The evaluation unit is operative to execute the aforementioned method by providing suitable electronic components such as integrated circuits, such as a microprocessor, an ASIC, and computer memory for storing computer readable code and data, such as EPROMs, EEPROMs, flash memory and so forth. In particular, the evaluation unit is operative to receive image data from the vehicle camera via the input connection, which comprises images frames.

25

30

The evaluation unit assigns transition values to pixels of the image frames using a transition indicator and assigns three-dimensional coordinates to the pixels using a model of the exterior scene. In particular, the model can be represented by computer readable instructions, which process input data from the camera or from the cameras and assign three-dimensional coordinates to the pixels of the image frames.

Furthermore, the evaluation unit is operative to use the assigned three-dimensional coordinates and movement data of the vehicle to assign trajectories to the pixels. The movement data, which is received over the input connection, is provided by one or more vehicle sensors, such as a speed sensor, a steering angle sensor or camera sensors. The input connection may comprise separate connection lines for camera data and other data. In particular, the velocity, the steering angle or the ego-motion data may be provided by an automotive data bus, such as a CAN bus.

Furthermore, the evaluation unit is operative to compute, for pixels to which a trajectory is assigned, a respective expectation value for the pixels, which is derived from a transition value of a starting pixel of the trajectory, and to compute a satisfied transition value for the pixel location. The satisfied transition value is derived from transition values along the trajectory that begins with the starting pixel for a set of frames in the circular buffer that are adjacent to the frame of the starting pixel.

Furthermore, the evaluation unit is operative to accumulate the expectation value and the satisfied transition value over

time and to derive a cleanliness value from a ratio of the accumulated satisfied transition value to the accumulated expectation value.

5 Specifically, the respective accumulated values may be stored in arrays $Sat(x, y)$ and $Tot(x, y)$, which may be stored in a permanent memory such that they are available when the evaluation unit is powered on. For faster access, the arrays may also be provided in a read and write memory during operation of the evaluation unit.

Furthermore, a results buffer can be provided for storing the cleanliness values $C(x, y)$ and an output connection can be provided for outputting the cleanliness values $C(x, y)$. However, the input connection may also be used as output connection in case of a two-way data traffic. In particular, the evaluation unit may comprise a connection for connecting it to a further image processing unit, which evaluates the image frames using the information about the lens soiling.

20

Furthermore, the evaluation unit may comprise a power connection for receiving electric power from a battery power supply of a car.

25 In particular, the buffer for storing the image frames can be provided as a circular buffer. The circular buffer may be realized using pointers to the memory areas of the respective frames. Thereby, a shifting of memory content can be avoided and the last frame may simply be overwritten by a new incoming frame.

30

According to a further embodiment, the evaluation unit is operative to compare the accumulated expectation value with a predetermined threshold and to set the cleanliness value of the pixel to the ratio of the accumulated satisfied transition value to the accumulated expectation value, if the expectation value is greater than the pre-determined threshold. The cleanliness value of the pixel is represented as "undetermined" if the expectation value is less than the pre-determined threshold.

10

According to a further embodiment, the evaluation unit is operative to compute the transition values as a two-dimensional convolution of a filter with pixels of an image frame pixel.

15 According to yet a further embodiment, the evaluation unit is operative to compare the pixel transition values with a detection threshold and to set the respective transition value to one if the respective transition value is equal or greater than the detection threshold.

20

According to yet a further embodiment, the evaluation unit is operative to divide the image frames into disjoint subregions and to compute low-resolution transition values by selecting maximum values of the transition values in the subregions.

25 Furthermore, the evaluation unit is operative to store the low-resolution transition values in the buffer for a predetermined number of image frames and to use the low-resolution transition values to compute the corresponding expectation values and the corresponding satisfied transition values.

30

According to yet a further embodiment, the evaluation unit is operative to receive and process movement data from a vehicle speed sensor and from a steering angle sensor.

- 5 According to yet a further embodiment, the evaluation unit is operative to receive and to process movement data at least one camera sensor, such as a camera sensor of a surround camera system.
- 10 According to yet a further embodiment, the evaluation unit is operative to assign motion vectors to the pixels of the image frames according to a simplified model of a surrounding scene, the simplified model comprising a horizontal ground surface, a first region of vertical objects and a second re-
- 15 gion of vertical objects.

According to yet a further embodiment, the evaluation unit is operative to identify sets of nearby frames for which the vehicle is moving essentially in a straight line, within a pre-

20 determined margin, and to update the accumulated expectation value and the accumulated satisfied transition value for the identified frames only.

According to a further aspect, the present specification discloses a kit with the aforementioned evaluation unit and at

25 least one vehicle camera. The vehicle camera is connectable to the evaluation unit, directly via a dedicated cable or indirectly via an automotive data bus. The kit may furthermore comprise a data transmission cable, the camera and the

30 evaluation unit are adapted for a data transmission from the camera to the evaluation unit, for example by providing a pre-defined data exchange protocol on both devices.

In a further aspect, the present application discloses a vehicle with the aforementioned kit, in which the evaluation unit is connected to an automotive data bus of the vehicle, at least one camera is mounted to the vehicle, and at least one camera is connected to the evaluation unit.

The subject matter of the present application is now explained with respect to the following Figures in which

10

Figure 1 shows a lens image of a soiled lens,

Figure 2 shows a lens image of a wet lens,

Figure 3 shows a flow diagram of a method according to the present application,

15 Figure 4 shows a simplified scene with camera motion vectors,

Figure 5 shows a first camera with a first scene taken by a vehicle camera,

Figure 6 shows a transition request diagram indicating a total number of transition requests per pixel for the scene of Fig. 5 after 10 frames,

20

Figure 7 shows a lens contamination diagram for the scene of Fig. 5 after 10 frames in which clean pixels are indicated,

25 Figure 8 shows a second camera image with a second scene taken by the vehicle camera of Fig. 5,

Figure 9 shows a transition request diagram indicating a total number of transition requests per pixel for the scene of Fig. 8 after 200 frames,

30 Figure 10 shows a lens contamination diagram indicating for the scene of Fig. 8 after 200 frames in which clean pixels and artifact pixels are indicated,

Figure 11 shows a third camera image with a third scene taken by the vehicle camera of Fig. 5,

Figure 12 shows a transition request diagram indicating a total number of transition requests per pixel for the scene of Fig. 11 after 400 frames,

Figure 13 shows a lens contamination diagram indicating for the scene of Fig. 11 after 400 frames in which clean pixels and artifact pixels are indicated,

Figure 14 shows a fourth camera image with a fourth scene taken by the vehicle camera of Fig. 5,

Figure 15 shows an transition request diagram indicating a total number of transition requests per pixel for the scene of Fig. 14 after 700 frames,

Figure 16 shows a lens contamination diagram indicating for the scene of Fig. 14 after 700 frames in which clean pixels and artifact pixels are indicated.

In the following description, details are provided to describe the embodiments of the present specification. It shall be apparent to one skilled in the art, however, that the embodiments may be practised without such details.

A common assumption in imaging systems is that the radiance emitted from a scene is observed directly at the sensor. However, there are often physical layers or media lying between the scene and the imaging system. For example, the lenses of vehicle cameras, consumer digital cameras, or the front windows of security cameras often accumulate various types of contaminants over time such as water, fingerprints, dust and dirt. Artifacts from a dirty camera are shown in Figs. 1 and 2.

These artifacts can be disturbing for users but more importantly, can seriously impair the analysis of the scene by automatic methods. For example, in the automotive area, the wrong analysis of a scene may lead to the malfunction of security system when they are needed in case of an accident; false alarms, or unwanted action from the vehicle like an erroneous automatic break action. These artifacts can cause potentially life-threatening situations.

Because many camera systems are automated and are not often inspected, an automatic way of detecting such artifacts is needed. A method that detects that there are issues on the lens can order the lens to be cleaned, notify a human that it needs attention or disable or notify the methods that follow that that particular part of the image is not usable.

There are few methods on the literature that deal with this issue. The method disclosed in "Removing Image Artifacts Due to Dirty Camera Lenses and Thin Occluders", by J. Gu, R.

Ramamoorthi, P.N. Belhumeur and S.K. Nayar, in ACM Transactions on Graphics (Proceedings of SIGGRAPH Asia), Dec. 2009, attempts to detect and correct the artifacts in image sequences but needs that the output of computing the mean image and mean image derivative magnitude over time are mostly constant valued image-sized arrays except where artifacts are located (see Fig. 6 b) and c) in the aforementioned paper). This means that these quantities are only usable after a very large amount of frames, i.e., a long time, and that the scenes have to be totally uncorrelated.

While the first condition imposes a long detection time, the latter one is typically not applicable at all, since the

scenes in most real-life scenarios have always about the same type of content, e.g., a road below and sky above for automotive applications. The authors themselves recognize that lighting is typically correlated and propose a solution with inferior quality.

The method in the reference "Removal of Image Artifacts Due to Sensor Dust" by C. Zhou and S. Lin, Association for Computing Machinery, Inc., June 2007, reduces the appearance of dust in an image by first formulating a model of artifact formation due to sensor dust and then using contextual information in the image and a color consistency constraint. This method has a very narrow application range, i.e., the detection of dust particles, and minimizes a non-convex function using a computationally intensive and often instable method.

Finally, because the media obstructing the visual path is located at the lens, it either obstructs light completely, as shown in Fig. 1, or it makes the information very blurry in the soiled regions. The latter occurs because the media is not in focus and also because elements like water droplets have a magnifying effect, as shown in Fig. 2. An approach which detects blurred regions works best if the entire image or very large portions of it are being considered, since it is very unlikely that an image taken from a realistic environment is blurred everywhere.

Such an approach does not work well for small regions because, although a soiled region implies very often that the region is blurred, a blurred region does not imply that it is soiled. The image can contain blurred regions simply because there is no texture on the scene at that location. That would

be the case of an image of a white featureless wall, the door of a vehicle, a clear sky, etc. So, there is an ambiguity in knowing whether a blurred/texture-less region in the image is caused by lack of texture in the scene or soiling on the lens. This ambiguity is solved by the proposed method.

By using an artifact detection method according to the present specification this ambiguity can be reduced or even resolved. According to the artifact detection method, areas with sharp edges can be marked as artifact free. This step is based on the consideration that a sharp edge on the image indicates that the lens is not soiled in that location.

According to a key aspect of the artifact detection method, a known movement of the vehicle is used to derive a trajectory of an edge point and to check if the edge point is observed at an expected location of the image. If an edge point is not detected at some location along its expected trajectory, the location corresponds to a location of the lens that is soiled and is marked as an artifact.

In a first step, the magnitude of derivatives, E , in the image, I , is computed for every frame. In particular, the derivative E can be obtained by a convolution with a filter that approximates a derivative. Absolute maximum and minimum values of the second derivative have been found particularly useful for edge detection. The sum of the second derivatives in all spatial directions is also referred to as a "Laplace operator".

According to the present method, the direction of the derivative can be disregarded and, therefore, a simple isotropic Laplacian filter can be used,

5

$$F = \begin{bmatrix} 1 & 1 & 1 \\ 1 & -8 & 1 \\ 1 & 1 & 1 \end{bmatrix}. \quad (1)$$

This approximation is isotropic, i.e., puts equal weight on vertical, horizontal and diagonal derivatives. Other approximations of the two-dimensional Laplace operator can also be
10 used, such as

$$F = (1 - \gamma) \begin{bmatrix} 0 & 1 & 0 \\ 1 & -4 & 1 \\ 0 & 1 & 0 \end{bmatrix} + \gamma \begin{bmatrix} 0.5 & 0 & 0.5 \\ 0 & -2 & 0 \\ 0.5 & 0 & 0.5 \end{bmatrix},$$

15 wherein γ is between 0 and 1. The derivative E is then computed as

$$E = |I * F|, \quad (2)$$

wherein $I * F$ stands for a two-dimensional convolution of the filter with the pixel values and " $||$ " stand for a magnitude
20 such as the L2 norm. The pixel values $I(x, y)$ represent an intensity, such as a grey scale intensity, an intensity in a colour channel or also an effective intensity of multiple colour channels.

25 In terms of pixel coordinates x, y the discrete two-dimensional convolution can be written out as

$$G[x,y] = \sum_{u=-k}^k \sum_{v=-k}^k F[u,v] I[x-u,y-v],$$

wherein $k = 1$ for the abovementioned filter and wherein the coordinate $(-1, -1)$ corresponds to the upper left corner of the matrix F .

- 5 Instead of a Laplace filter, other derivative operators such as the Sobel or Prewitt derivative operators can be used to identify changes within the images.

10 The values along the borders of the image frame can be computed by imposing special conditions, for example by creating an outer frame, wherein the pixels have the same or similar values as the outer pixels or zero values. According to another embodiment, the convolution is only computed for the inner pixels of the image frame.

15

The method can be further simplified in that the actual value of the derivative is disregarded. Instead, the occurrence of a transition is computed. In particular, a probability of finding a transition is computed as the normalized minimum
 20 between a detection threshold, $ThEdge$, and the derivative values, E ,

$$T_{HRes} = \min\left(1, \frac{E}{ThEdge}\right). \quad (3)$$

The detection threshold $ThEdge$ is chosen sufficiently high, such that only the clearly visible transitions are detected.

25

According to a further aspect of the present method, a lower resolution version of T_{HRes} is computed, wherein each pixel

corresponds to non-overlapping $N \times N$ pixels of the original resolution. By way of example, $N \times N$ can be chosen as 20×20 pixels. By using the lower resolution version, the computational load can be decreased and a constraint of high precision point trajectories can be removed.

A low resolution probability of detecting transitions in the image is computed as

$$T(x_0, y_0) = \max \{ T_{HRes}(x, y) : (x, y) \in \hat{U}(x_0, y_0) \}, \quad (4)$$

wherein each region $\hat{U}(x_0, y_0)$ contains $N \times N$ pixels. The low resolution probability is assigned to all pixels of the subregion. The pixels of the subregion correspond to an effective "N x N pixel" (x_0, y_0) , which is stored as a single low-resolution pixel in computer memory.

15

For example, a 1280×800 image frame can be divided into non-overlapping or disjoint quadratic subregions with sizes 10×10 , 16×16 , 20×20 or 40×40 . According to other partitions there may also be regions $\hat{U}(x_0, y_0)$ with different sizes. This may be the case if an image frame cannot be subdivided into $N \times N$ regions without remainder. Furthermore, the subregions may also comprise rectangular subregions, in particular near the boundary of the image frame.

20

In order to analyse the trajectories over time, the probabilities of detecting a transition for a few frames is stored in a computer memory. A method of determining the number of frames to be stored is described further below.

25

If there is a high probability for a transition at point (x_0, y_0) in frame f , the point is likely in a clean region of the lens. In this case, the transition point will be visible along its natural trajectory in the next frames unless the
 5 lens is soiled.

The conditional probability that the lens is clean in a point (x, y) along the trajectory of a transition point is represented as the conditional probability of finding a transition in
 10 that trajectory in a nearby frame with a frame index $\hat{f}$, under the condition that a transition was detected at point (x_0, y_0) . In mathematical terminology, this conditional probability is defined as :

$$C_f(x, y) = p(\exists_{\hat{f} \in [f+1, f+F]} T_{\hat{f}}(x, y) | T_f(x_0, y_0)) \quad (5)$$

15

In other words, the right hand side refers to the probability that there exists at least one frame, with index in $[f+1, f+F]$, such that it has a transition $T(x, y)$ for this frame index, under the condition that the transition in point (x_0, y_0) has value $T(x_0, y_0)$, and under the condition that (x, y)
 20 is a point on an estimated transition point trajectory starting in (x_0, y_0) . The transition values T can have any value between 0 and 1.

25 By applying Bayes law and using the previously defined approximation to the probability $T(x, y)$, the conditional probability can be computed as

$$C_f(x, y) = \frac{p(\exists_{\hat{f} \in [f+1, f+F]} T_{\hat{f}}(x, y))}{T_f(x_0, y_0)} = \frac{\max \{T_{\hat{f}}(x, y) : \hat{f} \in [f+1, f+F]\}}{T_f(x_0, y_0)}. \quad (6)$$

Similarly, given some probability that there is a transition at point (x_0, y_0) in frame f , it can be expected that this transition point is visible along a given trajectory in the previous frames as well. If f refers to the present frame, then
5 only previous frames are available.

Choosing a range of frames, in this case the frames between the index $f+1$ and $f+F$, takes into account that the velocity
10 estimate is not absolutely accurate and the velocity may change and therefore the time at which the transition will pass over the pixel (x, y) is not known exactly.

In general, the exact trajectory of the transition points in
15 the frames is only known exactly if the three-dimensional coordinate of the points and the camera motion are known. For this reason, a set of approximations is used that makes this problem tractable. Firstly, by using low resolution transition probability arrays, vectors with large deviations from
20 the assumed one often correspond to the same low resolution point, thus reducing the demand for high motion vector accuracy.

Secondly, $C_f(x, y)$ is computed for frames in which the vehicle
25 is moving in an approximate straight line. Thereby, the trajectory estimation is made simpler. The motion of the vehicle can be determined by using CAN data and the vehicle speed. This situation occurs most of the times and will in general lead to a valid approximation. However, all other motions are
30 also possible.

According to a third approximation, only those points are considered which are located on the floor, on a left vertical plane, or on a right vertical plane at a given constant distance from the vehicle. These last two approximations enable
5 the computation of approximate motion vectors, whose direction is essentially correct and which have only a small error in their size.

Fig. 4 illustrates an estimated three-dimensional surface and
10 vectors obtained from it. The scenario of Fig. 4 is such that the image processing system estimates a forward motion. Fig. 4 shows that some vectors are approximately zero. Since no trajectory is created in these regions, it is not possible to estimate soiling of the lens there. However, they occur
15 mostly in sky regions when the vehicle has forward motion and, therefore, are not so important. To determine soiling in all areas of the lens, other types of motion, such as curving to the left or right, are needed.

20 According to the model of Fig. 4, the scene is subdivided into a first region, which corresponds to the ground surface, a second region, which is left to a fixed point of the motion and a right region, which is right the fixed point. In the left and right region, the objects are assumed to have vertical
25 orientation.

An artifact detection according to the present specification is improved when the overall probability $\mathcal{C}(x,y)$ of the lens being clean is obtained over time and with a lot of accumulated
30 evidence. Let a transition request be the expectation that is created in point (x,y) from the fact that a transition was detected in point (x_0,y_0) in a nearby frame.

In this scenario, $C(x,y)$ can be computed as the ratio between the number of transition requests that are satisfied, $Sat(x,y)$, and the total amount of transition requests, $Tot(x,y)$, for that pixel. $Tot(x,y)$ and $Sat(x,y)$ are probabilities and may also be regarded as relative numbers of transition requests.

For each pixel (x, y) and each frame f , the following values are computed:

1. If it is determined that the probability of a transition is 1, or in other words, that $T(x_0, y_0) = 1$, it can be assumed that the lenses are clean in those pixels. For these pixels, the quantities

$$\begin{aligned} Sat_1(x_0, y_0) &= Sat(x_0, y_0)(1 - upEdge) \\ &\quad + upEdge \\ Tot_1(x_0, y_0) &= Tot(x_0, y_0)(1 - upEdge) \\ &\quad + upEdge \end{aligned} \quad (7)$$

are computed, where $upEdge$ is an Infinite Impulse Response (IIR) update constant with a large value. A value of $upEdge = 0.5$ is used to obtain the Figures 5 to 16 according to the experimental data. If $T(x_0, y_0) < 1$, the quantities are set to zero, $Sat_1(x, y) = Tot_1(x, y) = 0$, wherein (x_0, y_0) is equal to the pixel location (x, y) .

The arrays $Sat(x, y)$ and $Tot(x, y)$, on the right hand side of the equations refer to the values before the update step. By convention, update equations, such as the equations (7) to (12), represent instructions to read values from computer memory,

perform the computation on the right hand side of the equations and to store the resulting value in a memory area, which is represented by the left hand side of the equation.

- 5 The arrays $Sat(x,y)$ and $Tot(x,y)$ are initialized with zero values, so if $Sat(x_0,y_0)$ and $Tot(x_0,y_0)$ had not been updated previously the left hand values of equation (7) are equal to the constant "upEdge".
- 10 2. The quantities $Sat(x,y)$ and $Tot(x,y)$ are updated at coordinates (x,y) along the trajectory of a point (x_0,y_0) for which $0 < T(x_0,y_0) < 1$, to reflect how much of the transition detected in (x_0,y_0) is visible in a nearby frame at point (x,y) . The amount by which these quantities are updated is proportional
- 15 to the value of $T(x_0,y_0)$. Thus, a strong transition creates a strong expectation and a large update and a weak transition creates a weak expectation and a small update.

The update equation is applied two times in order to detect

20 if transitions of a frame $f - F$ in the past are detected in the subsequent frames $[f - F + 1, f]$, according to

$$\begin{aligned}
 Sat_2(x,y) &= Sat(x,y) \left(1 - upTraj \cdot T_{f-F}(x_0,y_0) \right) + upTraj \\
 &\quad \cdot \min\{T_{f-F}(x_0,y_0), \max\{T_{\hat{f}}(x,y) : \hat{f} \in [f - F + 1, f]\}\} \\
 Tot_2(x,y) &= Tot(x,y) \left(1 - upTraj \cdot T_{f-F}(x_0,y_0) \right) + upTraj \cdot T_{f-F}(x_0,y_0)
 \end{aligned} \tag{8}$$

and in order to detect if transitions in the current frame f

25 were detected in the previous frames $[f - F, f - 1]$, according to

$$Sat_3(x,y) = \tag{9}$$

$$Sat(x, y) \left(1 - upTraj \cdot T_f(x_0, y_0) \right) + upTraj \cdot \min\{T_f(x_0, y_0), \max\{T_{\hat{f}}(x, y): \hat{f} \in [f - F, f - 1]\}\}$$

$$Tot_3(x, y) = Tot(x, y)(1 - upTraj \cdot T_f(x_0, y_0)) + upTraj \cdot T_f(x_0, y_0)$$

where $upTraj$ is set to a small value. A value of $upTraj = 0.05$ was used to obtain the pictures 5 to 16 according to experimental data. If $T_f(x_0, y_0) = 0$ or 1, then the quantities Sat_3 and Tot_3 are set to zero, $Sat_3(x, y) = Tot_3(x, y) = 0$. Similarly, if $T_{f-F}(x_0, y_0) = 0$ or 1, then the quantities Sat_2 and Tot_2 are set to zero, $Sat_2(x, y) = Tot_2(x, y) = 0$.

The effective arrays Sat and Tot at each frame are then computed as the maximum of these quantities. Thus, for each pixel,

$$\begin{aligned} Sat(x, y) &= \max\{Sat_1(x, y), Sat_2(x, y), Sat_3(x, y)\} \\ Tot(x, y) &= \max\{Tot_1(x, y), Tot_2(x, y), Tot_3(x, y)\}, \end{aligned} \quad (10)$$

This computation is followed by a slight reduction step to enable change over time,

$$\begin{aligned} Sat(x, y) &= Sat(x, y)(1 - upTime) \\ Tot(x, y) &= Tot(x, y)(1 - upTime) \end{aligned} \quad (11)$$

where $upTime$ is a very small value, $upTime \ll 1$. A value of $upTime = 0.001$ was used to obtain the experimental results of Figs. 5 to 16. A final probability that the lens is clean at a particular point is then computed by

$$C(x, y) = \frac{Sat(x, y)}{Tot(x, y)}. \quad (12)$$

By representing the amount of transitions that were expected at point (x,y) , the value of $Tot(x,y)$ provides a measure of confidence for the value of $C(x,y)$. Smaller values of $Tot(x,y)$ indicate that few transitions were expected in that pixel and that a large importance is being given to the few observations that were made. Larger values of $Tot(x,y)$ imply that $C(x,y)$ is the result of a large set of accumulated data and, therefore, has a higher probability of being correct. This is in turn indicates a higher probability that the value is not affected by the abovementioned simplifying assumptions.

It is therefore recommended that the value of $C(x,y)$ is only taken into account if $Tot(x,y)$ is very close to 1. In the experimental results of Figs. 5 to 16, $C(x,y)$ was only considered for points in which $Tot(x,y) > 0.95$.

The number of frames F to use and store are the result of a compromise. On the one hand, F should be as large as possible, to enable the detection of very large blurry areas: the more frames F are used, the longer is the expected trajectory and deeper it can go into blurry areas.

On the other hand, storing many transition arrays uses up memory even though the transition arrays have low resolution. Furthermore, as the actual motion may change over time, choosing a larger value of F increases the differences between the actual motion vectors and the assumed ones.

In the experimental results of Figs. 5 to 16, the pre-determined constant F is set to $F = 10$ stored frames. This

means that a total of 11 frames are used. At a frame rate of 30 frames/sec this corresponds to 0.37 seconds, by which time the car has moved about 5.14 m, if it moves at a speed of 50 km/h.

5

Fig. 4 illustrates a method for detection of a lens contamination according to the present specification. In a first step 20, image frames of size 1280 x 800 pixels are received from a vehicle camera, which operates at a frame rate of 30 frames per second.

10

In a second step 21, a magnitude of a derivative is computed for the pixels of the image frames by computing the two-dimensional convolution of a filter that approximates a Laplacian operator.

15

In a third step 23, a transition probability is computed as the minimum of the derivative magnitude divided by a transition threshold "Th" and the value 1. The transition probability can also be regarded as the probability that the pixel corresponds to a feature point, such as an edge.

20

In a fourth step 24, the image frame is subdivided into disjoint 20 x 20 subregions and a low resolution transition probability of a 20 x 20 subregion is computed as the maximum of the previously computed transition probabilities over the respective 20 x 20 subregion. The low resolution probability is assigned to an effective pixel.

25

In a fifth step 25, pixels having a low resolution probability of 1 are marked as clear. In a sixth step 26, the set of low resolution transition values for $F + 1$ consecutive image

30

frames, which has been computed in the previous step 24 is stored in a circular buffer.

In a seventh step 27, the vehicle speed and steering angle are received from a CAN bus and in an eight step 28, estimated motion vectors are computed using the vehicle speed and steering angle of step 27 and a simplified model of the exterior scene.

According to the simplified model, the motion vectors are computed for regions of the scene, which are assumed to be non-moving and to have horizontal or vertical orientation. In particular, the road surface and artificial structures along the road side provide corresponding objects.

In a ninth step 29, the low resolution transition values of the $F + 1$ image frames in the circular buffer are used to find transitions along the estimated trajectories and to compare detected and expected transitions.

In a tenth step 30, the ratios of detected and expected transitions are accumulated, wherein the regions, which have already been marked as clean in the fifth step, are taken into account. The ratios are used in an eleventh step 31 to create a "cleanMap" which indicates locations of the lens which are clean, soiled or for which a contamination status cannot be derived with a sufficient certainty. In particular, the cleanMap can be represented by a memory area of a computer memory, which represents values that indicate which pixels correspond to clean areas of the lens and which pixels correspond to contaminated areas of the lens.

EXPERIMENTS

The image series of Figures 5 to 7 shows the results of the artifact detection method after processing 10 frames of a sequence of the front camera using a surround view system, in which the lenses have drops of water. The second image, Fig. 6, shows the requests that were detected at those pixels, with brighter regions being originated mostly by strong transitions, since they have a high impact on this map. The third image, Fig. 7, shows an estimate of the state of the lens. In this scenario, the white pixels mean that there isn't enough data to confirm any scenario, which is equivalent to $Tot(x,y) \leq 0.95$ for the threshold of Fig. 7. The grey pixels indicate that the lens is clean at those pixels.

The image series of Figures 8 to 10, 11 to 13 and 14 to 16 show the evolution of the estimates, as more measurements are collected. The second images, Figs. 9, 12 and 15 indicate that the amount of data that was expected for each pixel becomes progressively complete. At the same time, the third images, Figs. 10, 13 and 16, converge towards the estimate of the state of the lens.

In these experiments, it takes a long time before a complete and credible estimate of the lens soiling is formed, due to the lack of texture caused by motion blur and the wet road. In simpler scenarios, where there is a lot of texture, the estimate can be formed in very few frames already. In other words, it is possible to obtain a good estimate in only one or two update steps or time accumulation steps.

CONCLUSION

The experimental results of Figs. 5 to 16 demonstrate the effectiveness of a method in capturing the soiling status of the lens, both in locations that are densely populated with transitions and on locations that are not. Naturally, areas with little content take longer to detect.

According to a first embodiment, transitions are detected but not tracked. By not tracking transition points, the method results to be simple, yet powerful. In this context, a "tracking" comprises an explicit search or comparison to find out where a feature went on the subsequent frames. By contrast, the method according to the first embodiment is based on the assumption that the transitions will be detected somewhere along the expected trajectory. It does not require a confirmation that a transition at a point (x, y) actually corresponds to the transition that was detected at point (x_0, y_0) . It could actually be a transition from some other point (x_1, y_1) .

According to further embodiments, the accuracy of the method can be increased by tracking the detected transition with a matching method such as sum absolute difference (SAD), correlation, sum squared difference (SSD) or a scale invariant feature transform (SIFT). Thereby, a pixel location to which the transition has moved in a nearby frame can be predicted more accurately.

A more complex and less general second embodiment works in the following way:

1. Find feature points in an image, for example by using a Harris corner detector.
2. Determine the three-dimensional coordinate of that point
5 using a standard structure for the motion. Preferentially, the standard structure for the motion is provided by static objects, which do not have an independent motion.
- 10 3. Use the estimated 3D coordinate of each feature point and the movement of the vehicle, which is provided for example by using CAN data or explicit ego-motion computation, a trajectory that each feature point exhibits in the images in the next frames is determined. Similarly,
15 the same method can be applied to determine the trajectory exhibited in the images in the previous frames.
4. By using the trajectory of the point, the approach according the first embodiment can be applied directly: an
20 expectation value is computed that each feature point will be detected along a trajectory. This expectation is not met in locations where the lens are soiled. By measuring the ratio between expectations that are met and the total expectations for each point, a location where
25 the lenses are soiled can be detected.

According to a third embodiment the approach according to the second embodiment is applied to a stereo scenario:

- 30 1. Find feature points in an image using, for example, a Harris corner detector.

2. Determine the 3D coordinate of that point using stereo information.

3. Use the 3D coordinate of each feature point and the movement of the vehicle, which can be derived by using CAN data or explicit ego-motion computation, the trajectories of the feature points can be predicted for both of the stereo images in the next frames. The same approach can be applied to determine the trajectory exhibited in both images, in the previous frames.

4. By using the trajectories of the point, an expectation value is computed that each feature point will be detected along its respective trajectory and will not be met in locations where the lenses are soiled. By measuring the ratio between expectations that are met and the total expectations for each point, the locations of the lenses which are soiled can be inferred. These locations may in general be different for the left and right images of a stereo camera.

CLAIMS

1. A method for detecting contaminations on a lens of a vehicle camera, comprising
 - receiving image data from the vehicle camera, the image data comprising images frames,
 - assigning transition values to pixels of the image frames,
 - assigning three-dimensional coordinates to the pixels using a model of the exterior scene,
 - using the assigned three-dimensional coordinates and movement data of the vehicle to assign trajectories to the pixels, the movement data being provided by one or more vehicle sensors,for those pixels to which a trajectory is assigned:
 - computing an expectation value for the pixel, the expectation value being derived from a transition value of a starting pixel of the trajectory,
 - computing a satisfied transition value for the pixel location, the satisfied transition value being derived from transition values along the trajectory for a set of frames that are adjacent to the frame of the starting pixel,
 - accumulating the expectation value and the satisfied transition value over time,
 - deriving a cleanliness value from a ratio of the accumulated satisfied transition value to the accumulated expectation value.
2. A computer program product with a computer executable code for executing the method according to claim 1.

3. An evaluation unit, the evaluation unit comprising
- an input connection for receiving image data,
 - a buffer for storing a predetermined number of image frames, wherein the evaluation unit is operative to
 - receive image data from the vehicle camera via the input connection, the image data comprising images frames,
 - assign transition values to pixels of the image frames,
 - assign three-dimensional coordinates to the pixels using a model of the exterior scene,
 - use the assigned three-dimensional coordinates and movement data of the vehicle to assign trajectories to the pixels, the movement data being provided by one or more vehicle sensors,
- for those pixels to which a trajectory is assigned:
- compute an expectation value for the pixel, the expectation value being derived from a transition value of a starting pixel of the trajectory,
 - compute a satisfied transition value for the pixel location, the satisfied transition value being derived from transition values along the trajectory for a set of frames in the buffer that are adjacent to the frame of the starting pixel,
 - accumulate the expectation value and the satisfied transition value over time,
 - derive a cleanliness value from a ratio of the accumulated satisfied transition value to the accumulated expectation value.
4. The evaluation unit according to claim 3, wherein the buffer for storing the image frames is a circular buffer.

5. The evaluation unit according to claim 3 or 4, further comprising a first memory area for storing an array of accumulated expectation values and a second memory area for storing an array of satisfied transition values, wherein an element of the respective array corresponds to a pixel location.

6. The evaluation unit according to one of the claims 3 to 5, wherein the evaluation unit is operative to
- compare the accumulated expectation value with a pre-determined threshold,
to set the cleanliness value of the pixel to the ratio of the accumulated satisfied transition value to the accumulated expectation value, if the expectation value is greater than the pre-determined threshold
and to represent the cleanliness value of the pixel as "undetermined" if the expectation value is less than the pre-determined threshold.

7. The evaluation unit according to one of the claims 3 to 6, wherein the evaluation unit is operative to compute the transition values as a two-dimensional convolution of a filter with pixels of an image frame pixel.

8. The evaluation unit according to one of the claims 3 to 7, wherein the evaluation unit is operative to compare the pixel transition values with a detection threshold and to set the respective transition value to one if the respective transition value is equal or greater than the detection threshold.

9. The evaluation unit according to one of the claims 3 to 8, wherein the evaluation unit is operative to
- divide the image frames into disjoint subregions and compute low resolution transition values by selecting
 - 5 maximum values of the transition values in the subregions,
 - store the low-resolution transition values in the buffer for a predetermined number of image frames,
 - use the low-resolution transition values to compute
 - 10 the corresponding expectation values and the corresponding satisfied transition values.
10. The evaluation unit according to one of the claims 3 to 9, wherein the evaluation unit is operative to receive
- 15 and process movement data from a vehicle speed sensor and from a steering angle sensor.
11. The evaluation unit according to one of the claims 3 to 10, wherein the evaluation unit is operative to receive
- 20 and to process movement data at least one camera sensor.
12. The evaluation unit according to one of the claims 3 to 11, wherein the evaluation unit is operative to assign
- 25 motion vectors to the pixels of the image frames according to a simplified model of a surrounding scene, the simplified model comprising a horizontal ground surface, a first region of vertical objects and a second region of vertical objects.
- 30 13. The evaluation unit according to one of the claims 3 to 11, wherein the evaluation unit is operative to identify frames for which the vehicle is moving essentially in a

e straight line, and to update the expectation value and the satisfied transition value for the identified frames.

- 5 14. A kit, the kit comprising the evaluation unit according to claim 13, and at least one vehicle camera, the vehicle camera being connectable to the evaluation unit.
- 10 15. A vehicle comprising the kit of claim 14, wherein the evaluation unit is connected to an automotive data bus of the vehicle, the at least one camera is mounted to the vehicle, and the at least one camera being connected to the evaluation unit.



Fig. 1





Fig. 2






Fig. 4




Fig. 5




Fig. 6




Fig. 7




Fig. 8




Fig. 9




Fig. 10




Fig. 11




Fig. 12




Fig. 13




Fig. 14




Fig. 15




Fig. 16

INTERNATIONAL SEARCH REPORT

International application No
PCT/EP2016/063650

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

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, COMPENDEX, 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	EP 2 682 898 A1 (CLARION CO LTD [JP]) 8 January 2014 (2014-01-08) abstract; figures 2,14 paragraphs [0001], [0002], [0007], [0014] - [0016], [0020] - [0029], [0042] - [0047], [0052] - [0056], [0060] - [0076] -----	1-15
A	US 2014/321701 A1 (HALIMEH JAD [DE] ET AL) 30 October 2014 (2014-10-30) abstract paragraphs [0028] - [0033], [0080], [0081] -----	1-15
A	EP 2 682 315 A1 (CLARION CO LTD [JP]) 8 January 2014 (2014-01-08) abstract paragraphs [0012] - [0023], [0031], [0051] - [0053], [0061] -----	1-15



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

12 August 2016

Date of mailing of the international search report

25/08/2016

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

Eckert, Lars

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No

PCT/EP2016/063650

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
EP 2682898	A1	08-01-2014	CN 103529639 A 22-01-2014
			EP 2682898 A1 08-01-2014
			JP 2014030188 A 13-02-2014
			US 2014010408 A1 09-01-2014

US 2014321701	A1	30-10-2014	CN 103857567 A 11-06-2014
			DE 102012210116 A1 20-12-2012
			EP 2720912 A1 23-04-2014
			US 2014321701 A1 30-10-2014
			WO 2012171737 A1 20-12-2012

EP 2682315	A1	08-01-2014	CN 103533231 A 22-01-2014
			EP 2682315 A1 08-01-2014
			JP 2014011785 A 20-01-2014
			US 2014009616 A1 09-01-2014
