(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property Organization
International Bureau

(43) International Publication Date
11 June 2020 (11.06.2020)

WIPO | PCT

(10) International Publication Number
WO 2020/114941 A1

(51) International Patent Classification:
G06K 9/46 (2006.01) *G06K 9/34* (2006.01)
G06K 9/48 (2006.01) *G06K 9/32* (2006.01)
G06K 9/62 (2006.01) *G06N 3/04* (2006.01)
G06T 7/11 (2017.01) *G06N 3/08* (2006.01)

(21) International Application Number:
PCT/EP2019/083245

(22) International Filing Date:
02 December 2019 (02.12.2019)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
18210459.6 05 December 2018 (05.12.2018) EP

(71) Applicant: **AGFA NV** [BE/BE]; IP Department 3802 Septestraat 27, 2640 Mortsel (BE).

(72) Inventor: **VANDERSMISSEN, Eva**; c/o AGFA NV, IP-Department 3802 Septestraat 27, 2640 Mortsel (BE).

(74) Agent: **PIETERS, Dirk**; AGFA NV, IP Department 3802 Septestraat 27, 2640 Mortsel (BE).

(81) Designated States *(unless otherwise indicated, for every kind of national protection available)*: AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DJ, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IR, IS, JO, JP, KE, KG, KH, KN, KP, KR, KW, KZ, LA, LC, LK, LR, LS, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SA, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

(84) Designated States *(unless otherwise indicated, for every kind of regional protection available)*: ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, ST, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, KM, ML, MR, NE, SN, TD, TG).

(54) **Title:** IMPROVING SEGMENTATIONS OF A DEEP NEURAL NETWORK

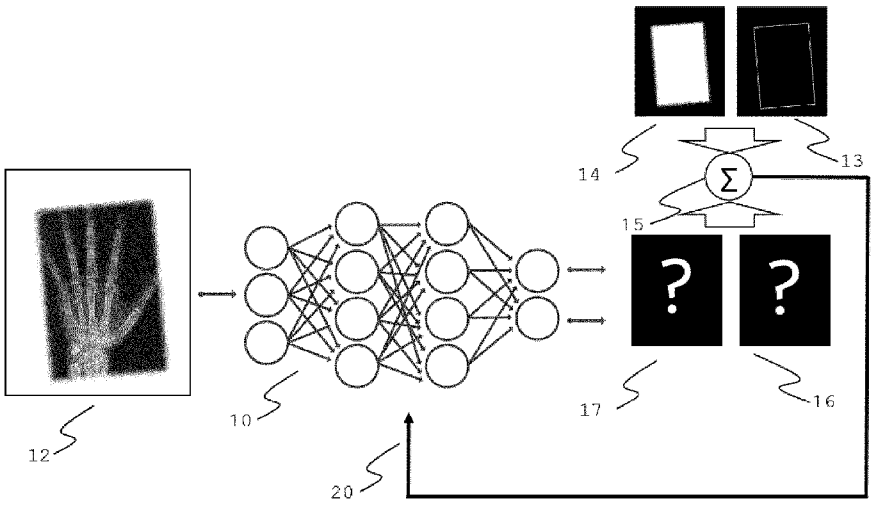


Fig. 1

(57) **Abstract:** This invention is related to a method to improve the performance of a deep neural network (10) for the identification of a segmentation target (111) in a medical image (12, 110), comprising the steps of performing n training steps on said deep neural network (10) for the identification of said region of interest on two different representations (13, 14) of the same segmentation target (111), said representations (13,14) being definitions of the same segmentation target (111).

WO 2020/114941 A1

[Continued on next page]

Published:
— *with international search report (Art. 21(3))*

IMPROVING SEGMENTATIONS OF A DEEP NEURAL NETWORK

Technical Field

[0001] The present invention relates generally to a method to improve the segmentation performance of a computer implemented deep learning algorithm, and this especially in the field of X-ray imaging. It is the purpose of this invention to generally improve the performance of particular segmentation tasks in radiography images, such as the detection of the collimation area, or the area in the image which is obscured by a body part, or the area representing bony structures.

Background of the invention

[0002] The performance of a medical imaging segmentation task, such as the determination of the collimation area in a radiography image is a typical problem in general radiography that up-to today finds many computer implemented approaches, which rely on different image processing methods. The collimation area is the exposed area of a detector forming the image that is produced by a collimated X-ray beam. Specifically for the collimation area detection problem, different conventional image processing solutions have been proposed in the art that rely at least on some form of edge detection. Collimation area detection was traditionally applied in practice by the radiographer on a freshly acquired image before it was sent to the radiologists reading station. The collimation step as such allowed the radiographer to remove superfluous areas of the acquired image without relevant information, while at the same time providing focus for the radiologist on the region of interest. Automation of this otherwise manual operation using image processing techniques reduces the workload for the radiographer.

[0003] Conventional image processing and feature detection techniques rely on known features of the elements within the image to be detected, such as for instance a sharp edge at the border of the collimation area, or

knowledge about the image information at the edges of an image. Also for instance, is it characteristic by nature for a collimated radiograph that the collimation area has a known rectangular shape, at least when the collimation area entirely fits the visible image.

[0004] Nevertheless, said conventional methods or task-specific algorithms often fail to correctly determine these collimation edges because of different reasons, such as for instance the high variability in the grayscale values on such an edge, which may be easily detectable by the human eye, but which may confuse conventional image processing algorithms when not properly implemented.

[0005] This results in the frequent occurrence of two types of error: either the collimation area is estimated as being too large (false negative errors), or else too much of the usable image is hidden by assumed collimation borders (false positive errors). The false positive errors are far less desirable, because in this case an undesired, but available, portion of the image is discarded (or cut away) which may lead to a retake of the image in the worst case scenario.

[0006] Relatively recently, computer based artificial intelligence is being considered as a potential solution for various image categorisation and segmentation problems. More specifically, artificial neural networks are trained (to progressively improve their ability) to do tasks by considering examples, generally without task-specific programming. For example, in image recognition, they might learn to identify images that contain cats by analysing example images that have been manually labelled as "cat" or "no cat" and using the analytic results to identify cats in other images. They have found most use in applications difficult to express with a traditional computer algorithm using rule-based programming.

[0007] Segmentation tasks in medical images fall in this category, and a lot of investigation is being done in this area on deep learning and in the application of deep neural networks. Deep learning computing are computing systems vaguely inspired by the biological neural networks that constitute animal brains. The neural network itself is not an algorithm, but rather a framework for many different machine learning algorithms to work

together and process complex data inputs. Such systems "learn" to perform tasks by considering examples, generally without being programmed with any task-specific rules.

[0008] A deep neural network is an artificial neural network with multiple layers which consists of interconnected nodes between the input and output layers. According to the input patterns, the deep neural network is presented with and using a learning rule, it can be trained to find a mathematical manipulation by modifying the weights of the connections between the layers to turn the input into the output. The network moves through the layers calculating a confidence level of each output. As such, a deep learning network may be trained to recognize the collimation area in a radiography image, and return the recognition result as a confidence level map. The user can review the results and compare them with the output the network should display (i.e. compare them with a confidence level map representing the ground truth, being the correct answer). Each mathematical manipulation as such is considered a node in a layer, and complex deep neural networks have many layers, hence the name "deep" networks. The goal is that eventually, the network will be trained to decompose an image into features, identify trends that exist across all samples and process new images by their similarities without requiring human input.

[0009] This principle may thus be applied to many different problem statements and is thus not limited to the application of finding the collimation area.

Summary of invention

[0010] The present invention provides a method for improving the performance of a deep neural network for the identification of a segmentation target (or a region of interest) in a medical image, comprising the steps of performing two training steps on said deep neural network for the identification of said region of interest on two different representations of the same segmentation target, said representations being a definition of a region and of a contour of said segmentation target, such as a detected

confidence level region map and a detected confidence level contour map, as set out in Claim 2.

[0011] More generally, the present invention provides for a method for improving the performance of a deep neural network on performing a segmentation task of identifying a region of interest in an input image which is a medical image, comprising the steps of, performing *n* training tasks on said deep neural network for the identification of said region of interest, wherein each training task is respectively performed on a different representation of said segmentation target, wherein $n \geq 2$,
performing a validation of one of said *n* representations of said segmentation target against the remaining of *n* representations, by means of applying shape similarity or shape matching algorithms between said one of said *n* representations and said remaining of *n* representations, returning a validated segmentation result comprising a validated representation of said segmentation target as a result for the identified region of interest.

[0012] Specific examples and preferred embodiments are set out in the dependent claims.

[0013] In the context of this invention, a detected confidence level region map represents the pixel wise probability of the presence of a part of the segmentation target in the medical image in the form of an image matrix. The probability of the presence of a pixel element at a pixel location is expressed as the pixel value. Similarly, a detected confidence level contour map represents the pixelwise probability of the presence of a part of a contour in the same medical image in the form of an image matrix.

[0014] The result achieved by the dual learning step as described above is that the deep neural network is trained twice on the same input data, resulting in a double output representing however the same result: the segmentation target represented by an area in the image (a detected confidence level region map), and the same segmentation target represented as a contour delineating the same area (a detected confidence level contour map). In the training step, the deep neural network is presented with two sets of ground truths for the two desired

outputs. Since the internal working of the neural network is different for the detection of the different representations of the segmentation result, it intrinsically means that there is truly supplementary information created that may be used to enhance the detection outcome.

[0015] Generally speaking, a deep neural network can thus be trained on different representations of the same segmentation result. In the context of this invention, a representation of a segmentation target may be a contour, a contour map, a probability contour map, a region map, a probability region map, a set of contour points, a set of corners delineating a polygon, a set of piecewise continuous functions or alike, that each may be embodied as an image mask representing said representations in the same or different spatial resolution as the original medical image or –alternatively- that can be described in another way (such as a list with coordinates of corners, or an alternative parameter space than the image space such as for example the polar representation of a line).

[0016] In the subsequent steps of the method of this invention, the detected confidence level contour map is used to identify a number of contour elements, which are tested and validated against the information that is present in the second output; the detected confidence level region map.

[0017] In the context of this invention, a contour map has to be understood as a representation of a contour in the form of a matrix (image matrix), a grid, a map or a mask. In this representation, the pixel value in a certain map location will determine the presence of a part of the contour at the location of that pixel. In case of a confidence level contour map, the pixel value in a certain map location will determine the probability of the presence of a part of the contour at the location of that pixel. Some locations in this confidence level contour map will have a higher probabilities associated to them than others.

[0018] In the context of this invention, the “detected confidence level contour map” is the contour map (in the form as explained above) that is calculated by the deep neural network from the provided input of a medical image.

[0019] The above mentioned principles may be applied to a segmentation algorithm that is for instance targeted toward the segmentation of the

collimation area in a radiographic image. The improved collimation segmentation algorithm of the invention will reduce the number of false positive and false negative errors, and thus will eventually reduce the risk for wrongly interpreting a part of the image when part of the region of interest is –for instance- mistakenly hidden. A false positive error in this context is the erroneous underestimation of the collimation area in a radiographic image. Another advantage is that consequently less manual intervention by a radiographer will be needed for making corrections to the collimation polygon. Also will the better detection performance of the exact collimation area lead to an improved image perception of the image when the collimation borders are more precisely black-bordered (i.e. visually collimated). Further advantages and embodiments of the present invention will become apparent from the following description and drawings.

[0020] The present invention can be implemented as a computer program product adapted to carry out the steps as set out in the description. The computer executable program code adapted to carry out the steps set out in the description can be stored on a computer readable medium.

Brief description of drawings

[0021] Fig. 1 is a graphic representation of the training process of a deep neural network [10]. The diagram represents the untrained deep neural network [10] which is fed with a set of medical images [12] as input for which a segmentation target (or ground truth) represented as a contour [13] and represented as a region [14] is available. A comparison [15] of the outputs [16] and [17] of the deep neural network (represented respectively by a contour and a region) is done with the respective ground truths [13] and [14]. A loss function is calculated indicating the relevance of the produced outputs [16] and [17], based on which an adjustment of the parameters [20] of the nodes in the deep neural network is performed. A new iteration step will then evaluate the performance of the adjusted deep learning network by means of the recalculation of the loss function.

[0022] Fig. 2 gives a schematic overview of the central position of the trained deep neural network [100] in a preferred embodiment of the invention, wherein a medical image [110] is presented as input to the deep neural

network, which produces two types of output: the confidence level contour map [120], and a confidence level region map [130]. The preferred embodiment relates to a segmentation method for identifying a collimation area in a radiography image.

[0023] Fig. 3 is a flowchart of a preferred embodiment of the invention that is used for segmenting a collimation area in a radiography image. [110] represents the original radiography image from which the collimation area [111] needs to be segmented. The radiography image [110] is presented to the deep neural network [100], which produces a confidence level contour map [120]. From the original radiography image [110], a safe initial contour [200] is estimated [201], which will be used as a starting point to compare and validate [305] the found improved proposal contour [210] against. The confidence level contour map [120] is subsequently and iteratively decomposed [300] (using a e.g. a RANSAC algorithm) into discrete contour elements [301]. This discrete contour element [301] is applied [304] to the safe initial contour [200] replacing one of the (former) contour elements, and as such changing into an improved proposal contour [210]. In a next step, a validation [305] of the changed contour element is performed by comparing the outputs of a cost function applied to both the safe initial contour [200] and the improved proposal contour [210]. The output of the cost function decreases in case there is a better measure of agreement d with the detected confidence level region map and/or the detected confidence level contour map. In case that the cost function of the improved proposal contour [210] is lower than the compared original contour [200], the contour element is validated and then proposal contour is propagated to the next iteration [305].

[0024] Fig. 4 is a flowchart depicting subsequent iterations of a preferred embodiment of the method of the invention. The flowchart represents the method that starts after the successful validation of a first contour element [301]. The data in the confidence level contour map [120] that contributed to the identification of the discrete contour element [301] is subsequently subtracted from the original confidence level contour map [120], resulting in another (more limited) confidence level contour map [121]. This limited

confidence level contour map [121] is then subsequently again further decomposed [300], resulting in a new discrete contour element [301']. This discrete contour element [301'] is then applied [304] to the previously improved proposal contour [210], resulting in a further improved proposal contour [210']. Again, a validation [305] is performed, propagating a positive result to the next iteration step for which again a (more limited) confidence level contour map [122] is constructed. At the end of the iterative process, a final contour result [910] is available.

Description of embodiments

[0025] In the following detailed description, reference is made in sufficient detail to the above referenced drawings, allowing those skilled in the art to practice the embodiments explained below.

[0026] As an initial step, in a preferred embodiment, one single deep neural network [10] is trained simultaneously on identifying two different representations [13, 14] of the same segmentation target [111], said representations being a definition of a region and of a contour of said segmentation target. In order to do so, a deep neural network defined by a number of nodes and a number of layers, is presented with a so-called training set of data [11]. This training set comprises a preferably high number of medical images [12] in which the segmentation target is visible, and for which the desired outputs [13, 14] are available. This means that for each of the training images, a description of the respective segmentation solutions have to be available. The segmentation solutions may be presented to the deep neural network in various formats, but both solutions need to be available; the segmentation target marked as a contour [13] and marked as a region [14]. The training set of medical images will, in practice, be prepared for training the deep neural network by means of a consistent set of digital images (characterized by a certain spatial resolution), and for instance a set of digital image masks representing the segmentation target as an area and as a contour. The digital mask representations may be provided in the same or other spatial resolution than the medical image itself. The set of digital image masks may be discrete maps (comprising only 1 or 0 as a presence indicator in

the respective pixel location), or by confidence level maps (providing a pixel-wise probability for the presence of a region of interest).

[0027] In another preferred embodiment, the single deep neural network may be replaced with two deep neural networks which will be trained separately on the different result types. One network will be trained to identify the segmentation target marked as a contour, while the other will be trained to identify the segmentation target marked as a region. This embodiment has the advantage that the overall performance of the pair of networks is higher in comparison with the performance of a single network, thanks to the specialisation of a network towards the performance of a single task. While this approach is expected to produce even better results (show better segmentation performance), this will come at the cost of more computing power.

[0028] In a preferred embodiment, the deep neural network is trained on the identification of the collimation area in a radiography image. The collimation area of a radiographic image is in fact the exposed portion of the image which is limited by the X-ray sources' collimators to reduce the exposed area of the patient to a minimum sufficient to perform the intended diagnose or reading. The collimation area typically has a rectangular shape, and contains the diagnostic data of the radiographic image. The unexposed area outside of this collimation area in this radiography image typically has a white color with little relevant information, as it is (usually completely) underexposed. Since collimation blades do not always block all the radiation, this white area partially still comprises image information. The result is that alongside the collimation area, there is a faint imprint of the surroundings of the collimated area, allowing sometimes a better orientation for the reader.

[0029] In general, there are multiple reasons why it is advantageous to identify the collimation area in a radiography: 1) it allows to focus image processing algorithms solely on the collimation area of this image (which most of the time improves the performance of such algorithms), 2) it makes it possible to selectively invert the color of the collimated area (i.e. the area outside the collimation) so that it is more comfortable for the

radiologist to read the image, as this operation reduces the contrast between the collimated area and the image itself.

[0030] In an alternative embodiment, the deep neural network is trained on the identification of more complex area's of interest, such as for instance the identification of bone structures in extremities or other body parts, the identification of the lungs in a chest, or for instance on the identification of malignant structures (tumors, lung nodules, …). It speaks for itself that the deep neural network will have to be trained in this case with the appropriate image data set, and that it is of essence that the targeted region of interest can be represented easily both as an area and as a contour.

[0031] After the training step of the deep neural network as explained above, a next step in the method of the invention is that a new medical image [110], for which the segmentation target has to be identified and located, is presented to the deep neural network in order to calculate the two outputs [120, 130] and where one of the outputs is the segmentation target that is represented as a detected confidence level contour map [120].

[0032] In a subsequent decomposition step [300], the detected confidence level contour map [120] will be decomposed into a discrete number of contour elements [301, 301'] that can be expressed in a (mathematically) analytical way, and which -when summed up- represent a completely closed contour of the targeted region of interest. In the case of the presence of multiple regions of interest in the medical image, the targeted region of interest may be expressed as a set of closed contours. In order to be able to describe such a proposal contour analytically, it may be decomposed into multiple contour elements, which are to be considered as a number of lines, curves or other shaped lines that may be described in an analytical way, and that together, when summed up, describe the entire proposal contour. The result of the sum of the found contour elements is a closed piecewise continuous function. It is important that the recomposed result is a closed contour, as this is a domain requirement for any segmentation problem. The decomposition step may be any type of algorithm that is capable of extracting certain contour features such as lines, arcs and

corners. Such a contour element may, in the simplest form, be a line segment that can be expressed easily in a mathematical and analytical way.

[0033] In a preferred embodiment, the decomposition algorithm is a RANSAC (Random Sample Consensus) estimation, that is an iterative method to estimate parameters of a mathematical model from a set of observed data that contains outliers (data points that are far away from the other points), and thus is less sensitive to noise. Alternative methods can be Hough transform, or Line Segment Detector algorithms.

[0034] The extracted contour elements from the decomposition step above are subsequently validated before they are accepted as a contributing contour element of the final segmentation result. In the context of this invention, this validation will be made against information in the output of the deep learning network which is not being used in the previous step for extracting the contour elements, and which is represented as an area (a detected confidence level region map). The validations thus take place at the level of the contour element, not yet at the recomposed contour.

[0035] In one embodiment of the invention, the validation of a certain contour element [301] is done by comparing a cost function [305] of a contour that is derived in a first instance from a so-called safe initial contour [200], with the same cost function when applied on a new contour proposal [210] in which the proposed contour element contributes . The new contour proposal [210] should have a lower cost function result in comparison with the previously proposed contour (which is the safe initial contour at the start of the iterative process) before the proposed contour element can be accepted as a valid contribution to the final solution. As a safe initial contour, a suitable contour should be chosen that falls within the boundaries of the medical image. Said safe initial contour should be chosen in a conservative fashion such that any risk to reduce this safe initial contour too much that it would exclude a part of the final solution is avoided. The most “conservative” safe initial contour that may be identified is the medical image boundary itself; if choosing this boundary as a first safe initial contour, then by definition there is no risk to exclude any data at

all. Performance of the algorithm may however be improved in case that the safe initial contour already is chosen to better approach the final result.

[0036] The initial contour may be estimated based on algorithms that are known in the art and that are capable of providing a relatively good initial estimate without risking to excluding any part of the segmentation target. Alternatively, it is also possible to estimate the initial contour based information about the applied positions of the collimator blades during the exposure by the X-ray modality.

[0037] The above mentioned cost function is based on a measure of agreement d of the tested contour with the detected confidence level region map and/or the detected confidence level contour map. This may be achieved for instance by calculating a weighted sum of how close the tested contour is to the detected confidence level contour map, and the Sørensen–Dice coefficient index of the inner part of the detected confidence level region map. The outcome of the cost function decreases for better candidates.

[0038] Additionally, other validation steps may be applied on the contour elements that depend on domain specific knowledge of the medical image. Validation steps (or acceptance criteria) relating to domain specific knowledge may accept or decline a certain proposed contour element as a valid solution in the case that certain criteria are met, or –respectively- not met. As an example that relates to the segmentation of a collimation area in a radiographic image, a possible validation criterion could be the fact that a typical collimation field has a quadrangular shape in case that it is entirely visible in the field of view, and that by consequence the contour elements are straight lines. Results that are not represented by straight lines may therefore be rejected. Also, contour elements that cut through a part of the detected confidence level region map may similarly be rejected.

[0039] In case that any of the validation steps fail on a proposed contour element during the decomposition step, and the contour element is consequently rejected, this contour element is replaced with the original corresponding contour element that was replaced by the newly proposed contour element in the previously proposed contour (which is the safe initial contour at the start of the iterative process). After a validation of a contour element, the

data of the confidence level contour map [120] contributing to the decomposition algorithm (e.g. the RANSAC estimation) of said contour element is removed from the confidence level contour map [120, respectively 121], eventually only partially removed since for some parts might contribute to other contour elements, resulting in a next (more data-limited) confidence level contour map [121, respectively 122] comprising less data than the upstream confidence level contour map. This allows the next step of the iteration to proceed with identifying the next contour element, until there is no further data present in the last confidence level contour map.

[0040] After a discrete number of iterations wherein contour elements are successfully identified, and validated or rejected, no further data will be available in the (more limited) confidence level contour map [122] which is the starting point of each next iteration step. At this point, all newly validated contour elements and all original corresponding contour elements that were not replaced by the newly proposed contour element due to their rejection, may be added up with each other to form a recomposed contour or the final contour result [910].

[0041] Even so, optional validation steps may be applied also on the level of the recomposed contour or the final contour result [910]. Such validation steps will also depend on certain characteristics for the segmentation result that is available as domain knowledge. For instance, bone segmentation should lead to a result wherein segmented area or areas clearly contain high electron density material, i.e. that the segmented area(s) are clearly coloured white in the radiograph. Or alternatively, for instance, domain specific knowledge could exclude certain proposed contours that have deviating shapes from the expected one or for instance seem to include multiple isolated areas, whereas this is not allowed.

[0042] As an example, the decomposition step for targeting a collimation area in a medical image will substantially differ from a decomposition step for targeting for instance lung tissue in a medical image. The domain knowledge about segmenting a collimation area may for instance allow assumptions with regards to the shape of the collimation area (which is

mostly rectangular, unless only partially visible in the field of view of the image), and also allows to make the consideration that a false negative result (collimation are is estimated too large) is less critical than a false positive result (where the collimation area would be estimated too small). Domain knowledge of the application will thus allow to implement intelligent decomposition decisions when determining the contour elements.

[0043] As an alternative embodiment, a similar approach may be conceived to start a decomposition step starting from the confidence level region map [130]. Based on the confidence level region map a region mask can be created (for example with a thresholding method). A further improvement of this mask can be created based on a so-called connected component analysis of the region mask and based on the neighbouring contour information each component has in the confidence level contour map. Components who are not connected to a minimal amount of contour map could be discarded, close components can be merged. Domain specific knowledge can be used to further improve the region mask. Such domain knowledge may be for instance a maximum number of components possible in one image or a definition of a minimum size component.

Claims

Claim 1. Method for improving the performance of a deep neural network [100] on performing a segmentation task of identifying a region of interest [111] in an input image [110] which is a medical image, comprising the steps of,

- performing two training tasks on said deep neural network for the identification of said region of interest on two different representations of said region of interest, said representations being a definition of a region and of a contour of said region of interest,
- obtaining a detected confidence level region map [130], representing a pixel wise probability of the presence of a part of the segmentation target in the medical image in the form of an image matrix, and a detected confidence level contour map [120], a pixelwise probability of the presence of a part of a contour in the same medical image in the form of an image matrix, for said input image [110] from said trained deep learning network,
- from said detected confidence level contour map, iteratively determine a set of discrete contour elements [301, 301'] to describe said detected confidence level contour map as a closed contour, by optimizing the sum of said continuous functions describing said discrete contour elements as a best fit to said detected confidence level contour map,
- performing a validation of each of said sum of said discrete contour elements by evaluating a measure d of agreement of an improved proposal contour [210, 210'] with said detected confidence level region map of said region of interest, said improved proposal contour [210, 210'] comprising said contour element,
- returning a final contour result [910] comprising said validated contour elements as a result for the identified region of interest.

Claim 2. Method according to Claim 1, wherein said two training steps comprise the following steps:

- performing a training step of said deep neural network [100] on identifying said region of interest, comprising the steps of:

+ presenting [201] a training data set of medical images [110] containing said region of interest [111] to said deep neural network [100] as input, said training data set comprising for each image an associated ground truth binary region

map [130] representing the location of said region of interest in said image,

- performing a training step of said deep neural network [100] on identifying said contour of interest, comprising the steps of:

+ presenting [201] a training data set of medical images [110] containing said region of interest [111] to said deep neural network [100] as input, said training data set comprising for each image an associated ground truth binary contour map [120] representing the location of a contour of said region of interest in said image.

Claim 3. Method according to any of the previous claims, wherein medical images are CR, DR or other radiography images,

Claim 4. Method according to any of the previous claims, wherein medical images are radiography images, and wherein said region of interest is the skeleton, are the lungs or bone structures

Claim 1.

Claim 5. Method according to Claims 1 to 4, wherein the described best fit is based on a RANSAC estimation or Hough transform.

Claim 6. Method according to Claims 1 to 4, wherein said contour elements are lines, curves or other continuous functions.

Claim 7. Method according to Claims 1 to 6, wherein said evaluation of said contour elements is performed by calculating a measure d' of agreement of said improved proposal contour [210'], said improved proposal contour [210'] comprising said contour element, with the detected confidence level region map [120], and comparing it with a measure d of agreement calculated for a previous proposal contour [210], said previous proposal contour [210] not comprising said contour element, with the detected confidence level region map [120], and accepting said contour element as a contour element of an improved proposal contour [210'] when $d'>d$,

Claim 8. Method according to Claims 1 to 7, wherein each of said contour elements is additionally validated against at least one rule that is defined by any criterion that may allow a contour element to be accepted/declined based on domain knowledge .

Claim 9. Method according to any of the previous claims, wherein said final contour result is validated against at least one rule that is defined by any

criterion that may allow a contour element to be accepted/declined based on domain knowledge.

Claim 10. A data processing system comprising means for carrying out the steps of the method of any of the previous claims.

Claim 11. A computer program comprising instructions which, when the program is executed by a computer, cause the computer to carry out the steps of Claims 1-10.

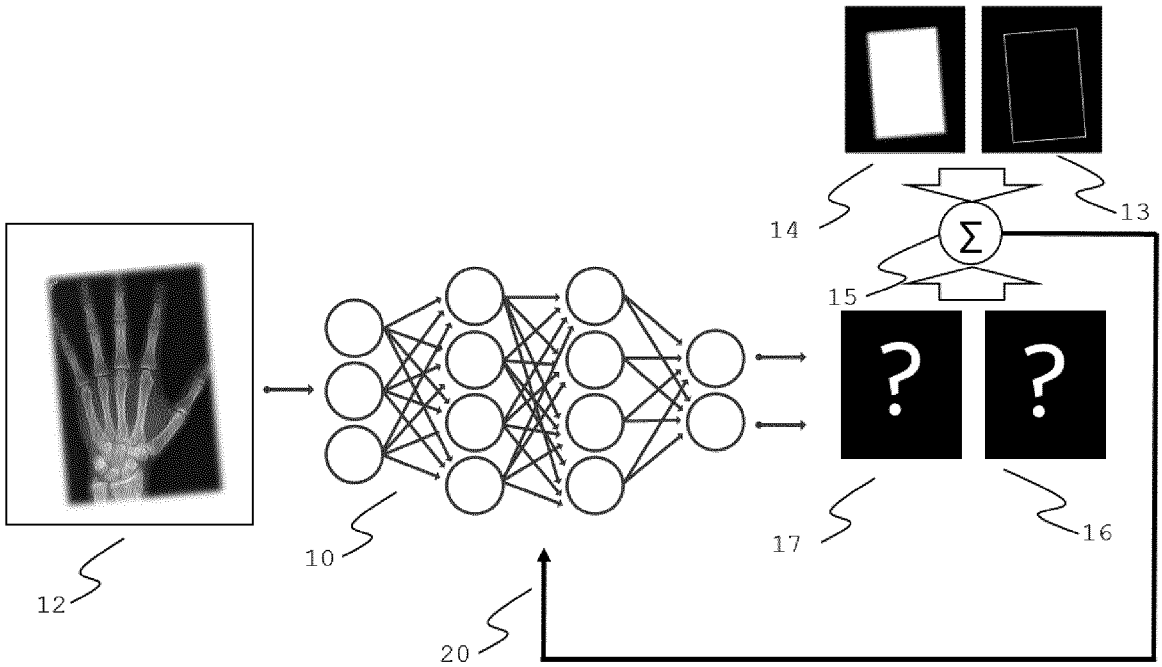


Fig. 1

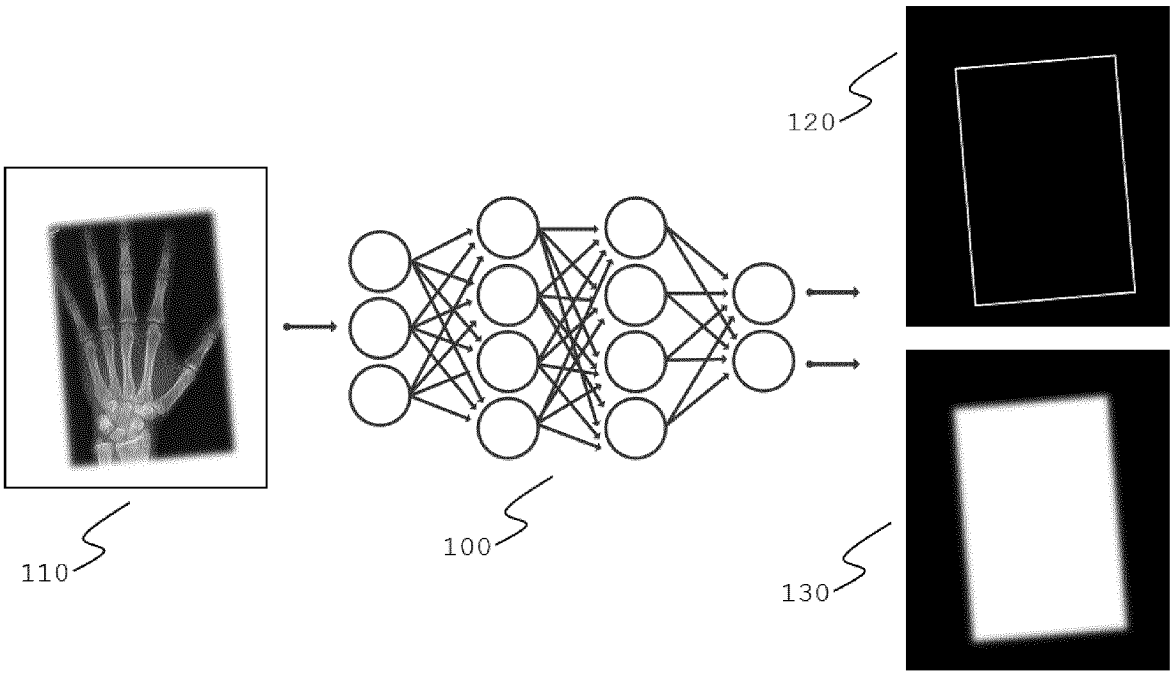


Fig. 2

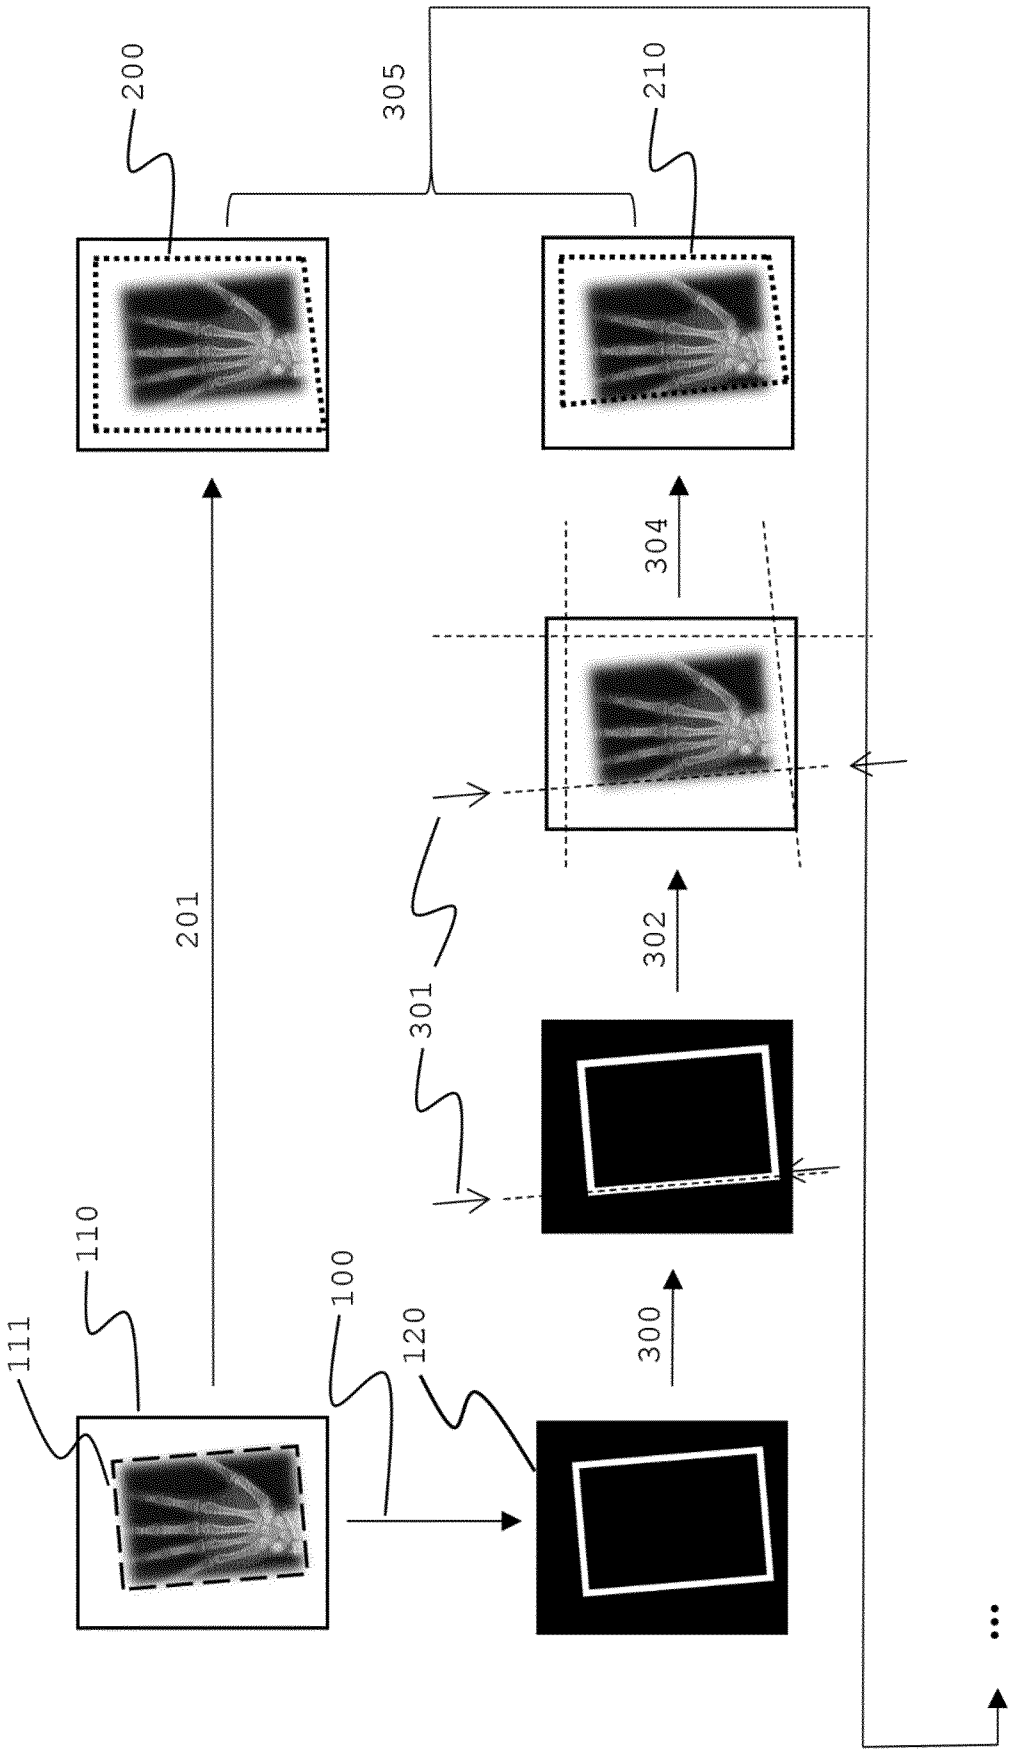


Fig. 3

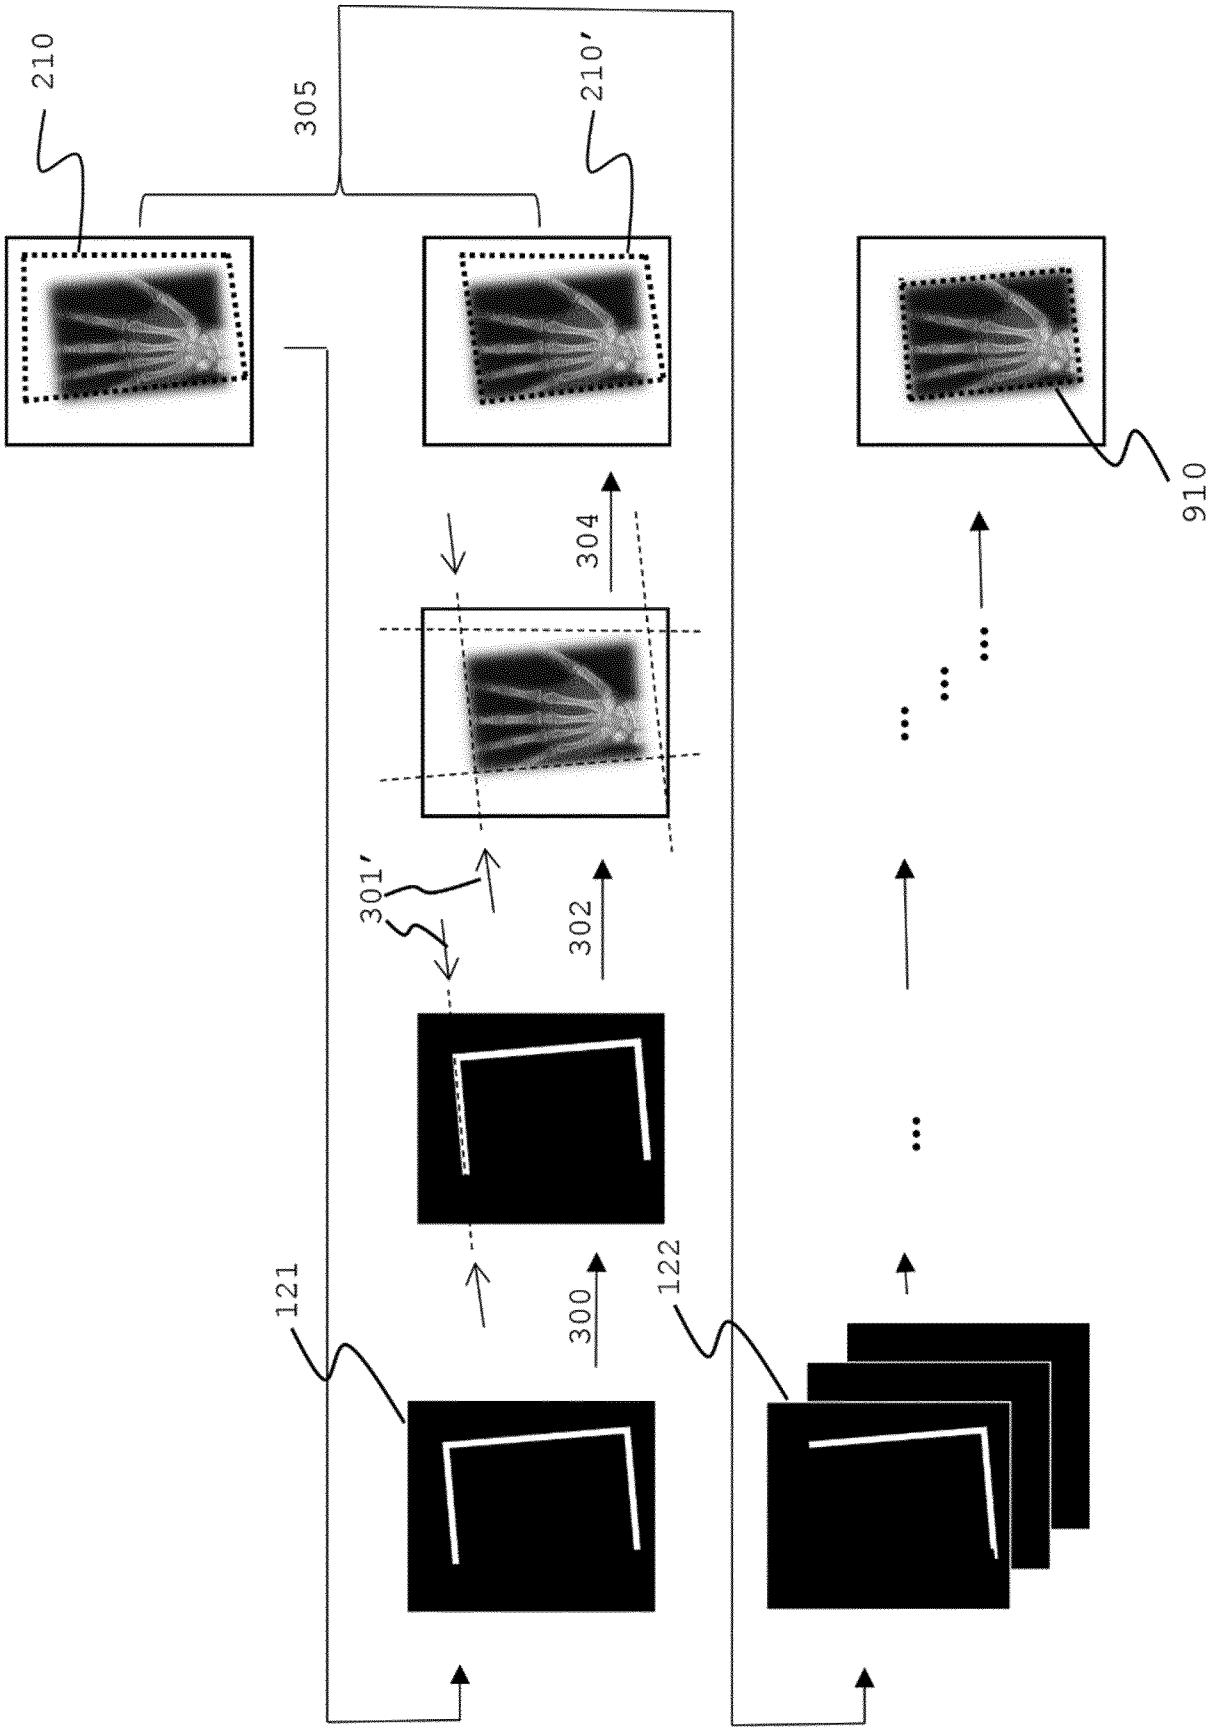


Fig. 4

INTERNATIONAL SEARCH REPORT

International application No
PCT/EP2019/083245

A. CLASSIFICATION OF SUBJECT MATTER
INV. G06K9/46 G06K9/48 G06K9/62 G06T7/11 G06K9/34
G06K9/32 G06N3/04 G06N3/08
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 G06T G06N

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

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	US 2017/287137 A1 (LIN ZHE [US] ET AL) 5 October 2017 (2017-10-05) paragraphs [0032] - [0038] paragraphs [0049] - [0068] paragraphs [0085] - [0104] figures 1-9	1-4,6,7, 10,11
A	US 6 212 291 B1 (WANG XIAOHUI [US] ET AL) 3 April 2001 (2001-04-03) column 4, line 21 - column 6, line 44 column 16, line 26 - column 19, line 4 figure 10	5,8,9

-/--

[X] Further documents are listed in the continuation of Box C.

[X] 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	Date of mailing of the international search report
28 February 2020	10/03/2020

Name and mailing address of the ISA/	Authorized officer
European Patent Office, P.B. 5818 Patentlaan 2 NL - 2280 HV Rijswijk Tel. (+31-70) 340-2040, Fax: (+31-70) 340-3016	Hermes, Lothar

1

Form PCT/ISA/210 (second sheet) (April 2005)

INTERNATIONAL SEARCH REPORT

International application No
PCT/EP2019/083245

C(Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	CHEN LIANG-CHIEH ET AL: "Semantic Image Segmentation with Task-Specific Edge Detection Using CNNs and a Discriminatively Trained Domain Transform", 2016 IEEE CONFERENCE ON COMPUTER VISION AND PATTERN RECOGNITION (CVPR), IEEE, 27 June 2016 (2016-06-27), pages 4545-4554, XP033021645, DOI: 10.1109/CVPR.2016.492 page 4545 - page 4548	1,10,11
A	EP 1 256 907 A1 (AGFA GEVAERT [BE]) 13 November 2002 (2002-11-13) paragraphs [0058], [0059]	1,5,8-11

1

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No
PCT/EP2019/083245

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2017287137 A1	05-10-2017	US 2017287137 A1	05-10-2017
		US 2018232887 A1	16-08-2018
US 6212291 B1	03-04-2001	NONE	
EP 1256907 A1	13-11-2002	DE 60133089 T2	30-10-2008
		EP 1256907 A1	13-11-2002
		JP 4307012 B2	05-08-2009
		JP 2003000579 A	07-01-2003