



(51) International Patent Classification:

G06Q 10/00 (2012.01) G06Q 40/08 (2012.01)

(21) International Application Number:

PCT/GB2021/050010

(22) International Filing Date:

04 January 2021 (04.01.2021)

(25) Filing Language:

English

(26) Publication Language:

English

(30) Priority Data:

2000076.6	03 January 2020 (03.01.2020)	GB
2000077.4	03 January 2020 (03.01.2020)	GB
2007465.4	19 May 2020 (19.05.2020)	GB
2016723.5	21 October 2020 (21.10.2020)	GB
63/198,628	30 October 2020 (30.10.2020)	US
2017464.5	04 November 2020 (04.11.2020)	GB

(71) Applicant: **TRACTABLE LTD** [GB/GB]; WeWork, 207 Old Street, London, Greater London EC1V 9NR (GB).

(72) Inventors: **RANCA, Razvan**; Tractable Ltd, Wework, 107 Old Street, London, Greater London EC1V 9NR (GB). **HORSTMANN, Marcel**; Tracktable LTD, WeWork, 207 Old Street, London, Greater London, EC1V

9NR (GB). **MATTSSON, Bjorn**; Tractable Ltd, Wework, 107 Old Street, London, Greater London EC1V 9NRgb (GB). **OELLRICH, Janto**; Tractable Ltd, Wework, 107 Old Street, London, Greater London EC1V 9NR (GB). **TEH, Yih Kai**; Tractable Ltd, Wework, 107 Old Street, London, Greater London EC1V 9NR (GB). **CHATFIELD, Ken**; Tractable Ltd, Wework, 107 Old Street, London, Greater London EC1V 9NR (GB). **KIRSCHNER, Franziska**; Tractable Ltd, Wework, 107 Old Street, London, Greater London EC1V 9NR (GB). **AKTAS, Rusen**; Tractable Ltd, Wework, 107 Old Street, London, Greater London EC1V 9NR (GB). **DECAMP, Laurent**; Tractable Ltd, Wework, 107 Old Street, London, Greater London EC1V 9NR (GB). **AYEL, Mathieu**; Tracktable LTD, WeWork, 207 Old Street, London, Greater London, EC1V 9NR (GB). **PEYRE, Julia**; Tractable Ltd, Wework, 107 Old Street, London, Greater London EC1V 9NR (GB). **TRILL, Shaun**; Tractable Ltd, Wework, 107 Old Street, London, Greater London EC1V 9NR (GB). **VAN OOSTEROM, Crystal**; Tractable Ltd, Wework, 107 Old Street, London, Greater London EC1V 9NR (GB).

(54) Title: METHOD OF DETERMINING INCONSISTENT DAMAGE TO PARTS OF A VEHICLE

Injecting Structured Accident Data
into Visual Damage Model 3200

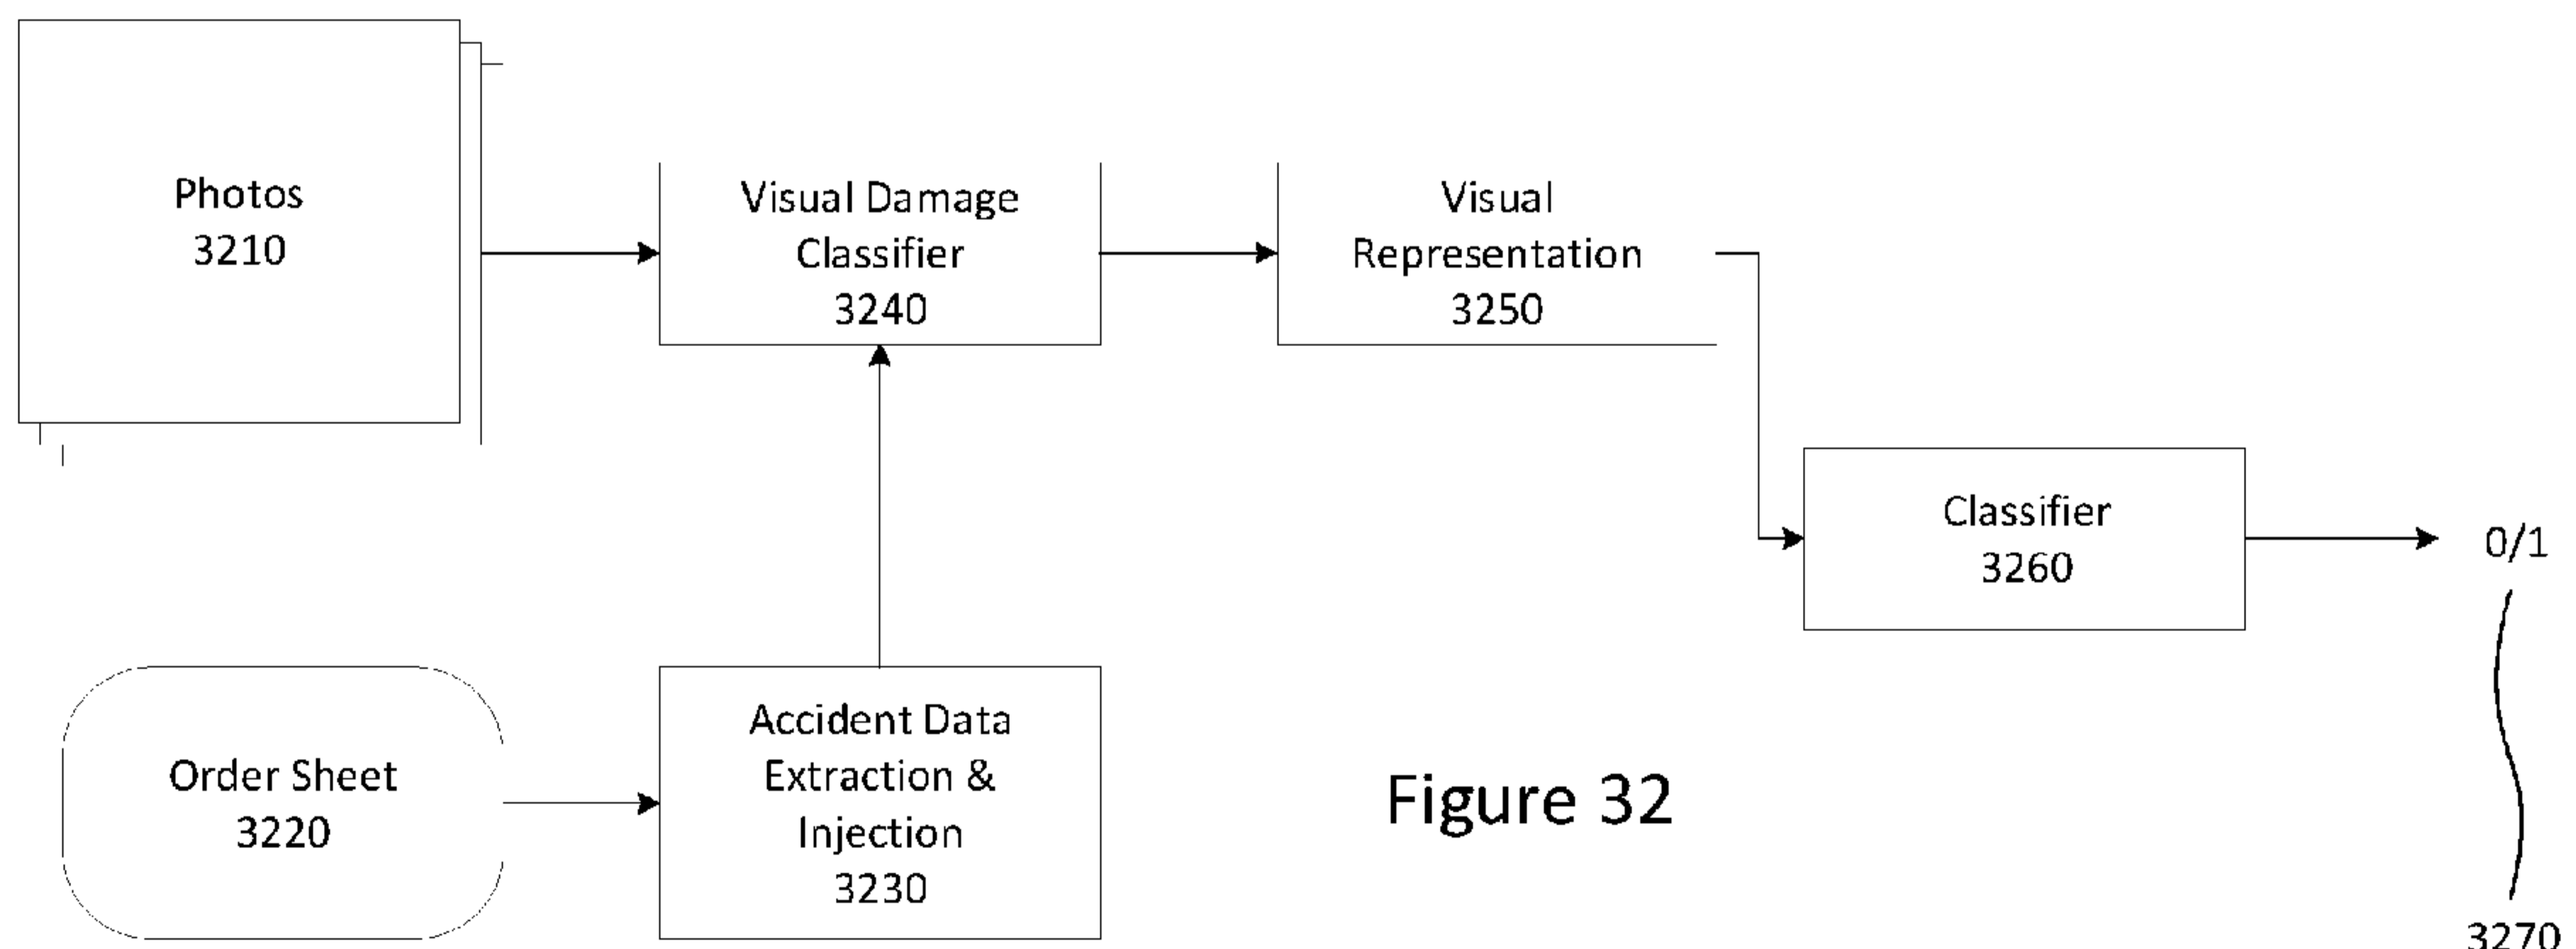


Figure 32

(57) Abstract: The present invention relates to the determination of damage to portions of a vehicle. More particularly, the present invention relates to determining whether determined damage to a vehicle is consistent with information provided as to the cause of the damage to the vehicle. Aspects and/or embodiments seek to provide a computer-implemented method for determining whether damage to a vehicle, which is determined using images of the damage to the vehicle, is consistent with information documenting the cause of the damage to the vehicle, for example insurance claim data or repair shop proposed repair data.

WO 2021/136941 A1

(74) **Agent: BARNES, Philip Michael**; Barnes IP Limited, Castle Hill House, 12 Castle Hill, Windsor, Berkshire SL4 1PD (GB).

(81) **Designated States** (*unless otherwise indicated, for every kind of national protection available*): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DJ, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IR, IS, IT, 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, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, WS, 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))*

METHOD OF DETERMINING INCONSISTENT DAMAGE TO PARTS OF A VEHICLE

Field

5 The present invention relates to the determination of damage to portions of a vehicle. More particularly, the present invention relates to determining whether determined damage to a vehicle is consistent with information provided as to the cause of the damage to the vehicle.

Background

10

Typically, and as shown in Figure 1, when a vehicle is involved in an accident (or is damaged) 105, the vehicle or its driver will be insured, and the driver will contact the relevant insurance company 110 to make a claim following a typical claim procedure 100.

15

The insurance company's estimation team 135 will then need to assess the damage to the vehicle and approve any claim, and the driver or insurer will then arrange for the vehicle to be repaired 145. Alternatively, the insurance company may make a cash settlement 150 in place of arranging or paying for repairs or may make a decision that the vehicle is a total loss 140 and compensate the insured party accordingly or arrange for a replacement vehicle to be procured.

20

As shown in Figure 1, the claim procedure 100 following an accident 105 requires the driver or insured party to call their insurer 110, and personnel at the insurer will follow a script 115 to receive and process the claim.

25

As part of the script 115, the insurer will obtain from the driver or insured party some information about the accident 105. Typically, the insurer will be provided with information about the insured person 120 (which may also include details of the vehicle and its condition etc that are provided during the call, or which are stored in the insurer's database and retrieved following receipt of the details of the insured person); details of the crash or accident 125, for example the circumstances and extent of the damage; and photos of the damage 130.

30

The photos of the damage 130 are typically taken by the driver or the insured party and can be of varying quality and comprehensiveness. Typically, photos are taken using phones equipped with cameras. Various problems can arise from this approach, including that too few photos are taken and provided to the insurer. Also, the photos taken may not be sufficiently well composed or may be of low quality due to the quality of the camera used to take the photos or the skill of the user.

35

The photos of the damage 130 can be provided to the insurer either via e-mail, facsimile or post, for example. This means there is typically a delay in the receipt of the photos 130 by the insurer, thus delaying the processing of the claim by the insurer and slowing down

the decision-making process as to whether the damage is a total loss 140, or whether a cash settlement 150 can be offered, or whether to arrange or allow the driver or insured part to arrange for repairs to the vehicle 145.

As part of the claim procedure, and more specifically the claim review procedure which
5 is carried out by the insurer to verify the costs of the proposed repair work by manually
assessing data provided by the client and any proposed repairer, the insurer may request
further information or claim data to be provided from the driver or insured party regarding the
accident. This may include details of the vehicle and its condition prior to any damage etc.
These are typically provided during a telephone call or are obtained having been stored in the
10 insurer's database, but sometimes requires the insurer to contact the insured party in a follow
up telephone call, letter or e-mail requesting the further details. Further, the insurer will require
sufficient details of the accident to be provided, along with sufficient photographs of the
damage for example, so this must be obtained during the first and any subsequent contact
with the insured party. The process of obtaining sufficient information can be slow, especially
15 if further requests for information are made in separate subsequent contacts with the insured
party, and thus can significantly delay the processing of an insurance claim. Further, the
proposed repairer may be required to send details of the proposed repairs, including for
example the labour tasks as well as any parts or materials costs, to the insurer for approval
prior to commencing work. The insurer can then assess whether the claim is covered by the
20 relevant policy under which the claim is made and determine whether the estimated costs of
repair can be verified and/or approved as may be appropriate.

Various tools and processes have been developed to assist vehicle repair businesses
and vehicle insurers respectively to prepare and approve repair proposals for damaged
vehicles, for example as a result of the vehicle being involved in an accident.

25 Vehicle repair businesses need to be able to itemise both the labour required and the
specific parts required in order to repair the vehicle, and then submit this for approval to an
insurer where the repair is covered by an insurance policy. Due to the large number of different
possible makes and models that might require repair, and the optional extras that might have
been fitted to the vehicle to be repaired, vehicle repair businesses typically have to use a
30 commercial database to identify the correct make, model, year of manufacture and options
fitted in order to correctly identify the parts that would need to be ordered if any need
replacement.

Insurers typically require vehicle repair businesses to submit evidence of the damage
to each vehicle and a detailed repair proposal that itemises the parts to be ordered and the
35 respective costs of each part along with detailed itemisation of the labour tasks and time that
will be required to carry out any repairs or replacement of parts. Preparing such detailed repair
proposals manually typically takes vehicle repair businesses a significant amount of time.

In different jurisdictions, different approaches are taken by both vehicle repair businesses (in respect of how repairs are carried out, what labour is deemed to be required, and preferences as to whether to repair or replace parts, for example) and insurers (in respect of what policies are applied when approving or rejecting proposed repairs, for example), depending on a variety of factors such as commercial pressures, regulation, consumer preference and typical insurance coverage. Thus, detailed repair proposals will differ between jurisdictions and what insurers are prepared to approve in a detailed repair proposal will also differ between jurisdictions.

Insurers, however, typically perform manual reviews on proposed repairs that are submitted for approval by vehicle repair businesses. As a result, the manual review process either requires a large workforce to perform the task of reviewing each submitted repair proposal or becomes a bottleneck in the repair approval process. For vehicle repair businesses, manual review can result in several disadvantages including delay in being able to begin repair work; further delays if the repair proposal is rejected by the insurer; and having to store customer vehicles for longer periods than necessary resulting in both higher storage space requirements and a higher probability of dissatisfied customers.

Across all jurisdictions, a variety of the above-described problems can result from manual preparation of proposed vehicle repairs and manual review of the proposed vehicle repairs by insurers.

Improvements to the claim procedure would enable repairs to be completed sooner and for insurers to reach decisions faster and more efficiently.

Summary of Invention

Aspects and/or embodiments seek to provide a computer-implemented method for determining whether damage to a vehicle, which is determined using images of the damage to the vehicle, is consistent with information documenting the cause of the damage to the vehicle, for example insurance claim data or repair shop proposed repair data.

According to a first aspect, there is provided a computer-implemented method for determining a measure of consistency between one or more damage indications to one or more parts of a vehicle, comprising the steps of: receiving a plurality of images of the vehicle; determining, per part, one or more classifications for each of the plurality of images using one or more trained models, wherein each classification comprises at least one indication of damage to at least one part; determining a measure of consistency between the one or more indications of damage; and outputting the measure of consistency of each of the parts of the vehicle wherein the indication of damage to at least two of the parts is determined to be correlated.

Providing images of damage to a vehicle and processing these to determine the extent of the damage to the vehicle and then correlating the determined extent of the damage to the vehicle with information pertaining to the cause of the damage to the vehicle can allow for any discrepancies between the damage to the vehicle and the cause of the damage to the vehicle to be identified.

Optionally, determining whether the at least one indication of damage to at least one part has a measure of consistency above a predetermined threshold that indicates the at least one indication of damage was caused in a single accident and/or incident.

Using a predetermined threshold can provide an objective method of determining the otherwise subjective measure of whether the cause of the damage to the vehicle is consistent with the damage to the vehicle determined from the images of the damage to the vehicle and/or can assist with identifying what damage to the vehicle, as determined from the images of damage to the vehicle, was caused by the indicated cause of damage to the vehicle.

Optionally, determination of a single accident and/or incident comprises any one or any combination of: location of the one or more indications of damage on the vehicle; type of damage; severity of damage; and/or one or more additional structured queries comprising point and/or of impact; vehicle motion state and/or manually generated text data.

The information pertaining to the cause of damage to the vehicle can comprise a variety of information about the cause(s) and/or damage.

Optionally, there is further performed the step of receiving accident data, wherein the accident data comprises data about the cause of the damage to the vehicle.

Accident data can be a rich source of information about the cause(s) of the damage to the vehicle.

Optionally, there is further performed the steps of processing the accident data to predict damage to one or more parts of the vehicle; and determining a correlation between the determined at least one indications of damage and the predicted damage.

Predicting the damage to the vehicle from the information about the cause of the damage to the vehicle can provide a set of data against which to compare the determined damage from the images of the damage of the data, which can in turn then be used to determine inconsistencies between the cause of the damage and the visible damage by comparing the predicted damage from the accident data and the determined damage from the image data.

Optionally, the output further comprises one or more indications of damage and/or relevant images corresponding to one or more vehicle parts determined to have a measure of consistency above a predetermined threshold.

Determining what damage is deemed to be consistent can allow for an output of the details of the damage deemed to be consistent, for example to allow repairs to be made for just the damage consistent with the cause of the damage.

Optionally, processing the accident data comprises the use of any or any combination of: Natural Language Processing; a neural network; a convolutional neural network; a recurrent neural network; and/or one or more trained classifiers.

A variety of methods can be used to process accident data in order to allow it to be used by the trained models/networks to determine predicted damage from the accident data for example.

Optionally, there is further performed the step of determining a plurality of parts of the vehicle that are represented in the plurality of images of the vehicle comprising the use of a plurality of classifiers. Optionally, each one of the plurality of classifiers is operable to detect each of the parts of the vehicle. Optionally, there is further performed the step of determining one or more sets of relevant images of each of the plurality of parts from the plurality of images of the vehicle. Optionally, there is further performed the step of determining a damage state of each of the parts of the vehicle using the one or more relevant images of each of the plurality of parts of the vehicle, the one or more damage states being determined as one or more quantitative values.

Determining which parts are present in the images of the damaged vehicle can allow for a damage determination for each part to be performed using all images in which that part is shown. Optionally, a classifier is trained for each part and can be used to determine which images show the part for which each classifier is trained.

Optionally, determining the one or more classifications and/or determining the one or more damage states comprises using a multi-instance learned approach. Optionally, determining the one or more classifications and/or determining the one or more damage states comprises using a graph neural network. Optionally, determining the one or more classifications for each of the plurality of images further comprises determining one or more feature representations

By using a multi-instance model, multiple images can be assessed together per part for which a damage state is being determined, which can allow for a more accurate assessment to be made as more images can be used to determine a consolidated damage state output. A graph neural network can be used to process multiple images per part to determine a consolidated damage state per part across all of the images showing each part, and can determine a link and/or relationship and/or connection between nodes of its network to represent this link/relationship/connection.

According to a further aspect, there is provided a computer-implemented method of training a neural network for determining a measure of consistency between one or more

5 damage indications to one or more parts of a vehicle, including neurons, each neuron being associated with weights, wherein the method comprises: obtaining training inputs, the training inputs comprising images of vehicles, classifications of damage to vehicles and/or a damage consistency threshold; for each training input, selecting one or more neurons based on their respective probability; adjusting the weights of the selected neurons such that the selected neurons are substantially operable to identify one or more indications of damage; processing the training input with the neural network to generate a predicted output; adjusting the weights based on the predicted output.

10 Training a network and/or model to determine consistency between information pertaining to the cause of damage to a vehicle and determined damage to a vehicle as assessed using images of the damage to the vehicle can allow for any discrepancies between the damage to the vehicle and the cause of the damage to the vehicle to be identified.

Optionally, a plurality of neural networks are trained for each of the parts of a vehicle.

15 Training a network/model per part allows for each network/model to be training substantially optimally to determine damage and/or inconsistency in damage to each part based on images of damage to that part and information pertaining to the cause of damage to the vehicle.

Brief Description of Drawings

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

Figure 1 shows a traditional insurance claim process according to the prior art;

25 Figure 2 shows a near real-time insurance estimation process according to an embodiment;

Figure 3 shows the process of determining an estimate for repair cost of a damaged vehicle from photos of the damaged vehicle and vehicle information according to an embodiment;

30 Figure 4 shows a flowchart outlining a process of predicting whether the damage to the vehicle is a total loss according to an embodiment;

Figure 5 shows a flowchart outlining a process of estimating the repair cost of a damaged vehicle from photos of the damaged vehicle and vehicle information according to an embodiment;

35 Figure 6 shows a flowchart illustrating in one embodiment that each part has a dedicated classifier and that the photos for each part are provided to a set of models for each part to produce initial estimates/pooled scores per part per model and then that the results are concatenated into a damage signature vector according to an embodiment;

Figure 7 shows a flowchart outlining a process of determining whether paint blending is required according to an embodiment;

Figure 8 shows a flowchart outlining a process of determining whether a part is undamaged, in need of repair or in need of replacement according to an embodiment;

5 Figure 9 shows a flowchart of a process for creating an estimate of repairs from photos input to the estimation method according to an embodiment;

Figure 10 shows the structure of a neural network which can be used to process the input images and metadata according to an embodiment;

10 Figure 11 shows a typical sequence of an insurance claim for vehicle damage, where the damage to the vehicle is repaired;

Figure 12 shows the sequence of an insurance claim process including automated verification of proposed repair costs according to a described embodiment;

Figure 13 shows a more detailed overview of the automated verification of proposed repair costs used in the embodiment shown in Figure 12;

15 Figure 14 shows an overview of an automated assessment platform according to an embodiment;

Figure 15 shows an overview of a comparison analysis process that can be used in an embodiment;

20 Figure 16 shows an illustration of a comparison analysis process that can be used in an alternative embodiment;

Figure 17 show how a pre-checked vehicle repair claim can be prepared and approved according to an embodiment;

Figure 18 shows an illustration of a generalised method of verifying a claim input according to an example embodiment;

25 Figure 19 shows an illustration of an embodiment of an approach to normalise parts of vehicles to assist with the verification process of the embodiments described herein;

Figure 20 shows a more detailed process of preparing an estimated repair schedule and/or repair cost according to another example embodiment;

30 Figure 21 shows an alternative representation of normalised parts for a vehicle according to an embodiment;

Figure 22 shows a paint check process according to an embodiment;

Figure 23 shows different example rules for a variety of jurisdictions as to what repairs would be permitted and categorise these alongside the respective decision made by the automated system;

35 Figure 24 shows a model architecture using a gradient booster tree with multiple decision trees according to an embodiment;

Figure 25 shows an equation illustrating a method of calculating paint cost according to an embodiment;

Figure 26 shows an overview of a damage determination process according to an embodiment;

5 Figure 27 shows the use of cropping to preserve image resolution in a damage determination process from images according to an embodiment;

Figure 28 shows the use of cropping and segmentation in a damage determination process according to an embodiment;

10 Figure 29 shows an overview of a damage determination process using a multi-image learning model according to an embodiment;

Figure 30 show an overview of a damage determination process using a multi-image graph model according to an embodiment;

Figure 31 shows the process of using a language model to extract information relevant to a determination of visual damage from images of the damage according to an embodiment;

15 Figure 32 shows the process of injecting structure accident data into a visual damage determination model according to an embodiment;

Figure 33 shows a method of predicting damage to grouped auxiliary parts of a vehicle where the grouped auxiliary parts are associated with a normalised part/panel of the vehicle according to an embodiment;

20 Figure 34 shows an example multi-task architecture according to an embodiment;

Figure 35 shows a method of jointly training across variably sized datasets using domain confusion loss according to an embodiment;

25 Figure 36 shows a method of jointly training across variably sized datasets using a subset of the larger dataset representative of the most comparable data according to an embodiment; and

Figure 37 shows an example domain confusion loss architecture usable with embodiments.

Specific Description

30

Referring to Figure 2, an example embodiment within a process for an insurer to handle an insurance claim for an accident will now be described.

35 In the following embodiments, and in the aspects and/or embodiments described above, the term “part” can be used interchangeably to refer to a specific part of a vehicle (i.e. a specific part of a specific make, model and year of vehicle) or a generic part of a vehicle as well as to a “normalised” or “standardised” part generic to most and/or all vehicles or to a region or zone of a generic vehicle or class of vehicle. Further details are provided in relation

to embodiments of this concept in Figures 19 and 21 and in the associated description for these figures set out below.

The insurance claim process is outlined in a flowchart 200 and begins with an accident 105 (or damage being discovered) involving a vehicle that is insured.

5 The party that is insured, or a party contacting an insurer on behalf of an insured entity or in respect of an insured vehicle, will then contact the insurer 110. The representative taking the telephone call on behalf of the insurer will run through a standard set of questions, or a decision tree of questions, to obtain a standard required set of information about the accident 120, 125, 130 (e.g. the name of the insured person/vehicle; details of the vehicle including make, model and condition; details of any insurance policy; standard security questions; 10 details of what happened to cause the accident; etc).

The phone call 110 and the party phoning the insurer providing the information requested by the insurer representative following the script 115 will mean that details of the insured person 120 (or insured party/vehicle) will be provided to the insurer as well as crash 15 details 125 providing information about the damage/accident.

During the call, the representative will ask the person phoning the insurer to provide photographs of the accident (and specifically photos of the damaged vehicle) 130 using the web app 205. If the person does not already have access to the web app 205, this can be provided by the representative by reading out a shortened URL or by sending a text message 20 or email or other notification to a device of the person phoning the insurer. Once the web app 205 has been loaded on the person's device, it can be used to take photographs using the device imaging function (e.g. a camera in a smart phone) to provide photos 130 to the insurer. The photos 130 will appear to the representative on their computer display shortly after these arrive at the insurer computer systems, and the representative can then guide the person on 25 the 'phone to take further photos.

In alternative embodiments, other means of providing photos to the insurer in parallel to the communication with the insurer can be used, so for example a chat interface can be used to communicate between the insured party and the insurer and this can be adapted to allow photographs to be taken and/or uploaded via the chat interface.

30 In this embodiment, the insurer provides a web application 205 that can be used by the person phoning the insurer 110 by accessing the web application 205 using a URL in a standard browser and taking photos 130 which are uploaded via the web application 205 and transmitted via the Internet to the insurer.

As the details of the insured person 120, crash details 125 and photos of the damage 35 130 are provided by the party phoning the insurer representative, the computer system used by the representative is configured to display a dynamic live estimate of the repair cost of the damage to the vehicle which is updated periodically as each item of new information is

received and processed by the insurer computer system (either by the representative entering it into the computer system as they follow the script for such calls or by receiving information from the insurer database about the relevant policy or by receiving photos of the damage via the web application 205).

5 The insurer representative can see a live damage estimate 210 displayed on the screen, which is determined using the method described in the following embodiments based on the provided accident/damage photos 130, and in some embodiments may also be presented with a level of certainty for that estimate 210 and whether any further photos of any damage are required to increase the level of certainty in the estimate. In some embodiments,
10 a required/threshold level of certainty is required in order for the representative to authorise one or more of the possible solutions for the insured party (i.e. total loss 140, repair 145 or a cash settlement 150).

Once there is an acceptable level of certainty in the estimate 210, or the representative has concluded their script 115, the three options typically available to the representative are
15 able to be chosen or offered by the representative if permitted (i.e. total loss 140, repair 145 or a cash settlement 150).

On the representative display, a total loss calculation 140 is made by a trained model which indicates whether, based on the information 120, 125, 130 provided by the party calling the representative 110, it appears that the value of the likely repairs to the vehicle indicate that
20 the vehicle is a total loss. The output of this total loss calculation 140 is a yes/no indication to the representative. The output can be presented to the representative while they are still on the telephone call with the person phoning the insurer, and the verification of the calculation can be performed by the representative or other insurer personnel during the telephone call. In some embodiments, an initial decision can be provided over the phone call but confirmation
25 will be made in some or all such decisions by insurance loss adjusters.

On the representative display, a cash settlement calculation 150 is made by a trained computer model and displayed to the representative which indicates whether, on the information 120, 125, 130 provided by the party calling 110 the representative, an amount has been determined according to the rules of the insurer that can be offered to the party calling
30 the representative as a cash settlement amount. The output of this cash settlement calculation 150 is a monetary amount that can be offered by the representative for immediate payment in settlement of the claim in lieu of arranging and paying for repairs to the vehicle. The output can be presented to the representative while they are still on the telephone call with the person phoning the insurer, and the verification of the calculation can be performed by the
35 representative or other insurer personnel during the telephone call, and the offer can be made by the representative or other insurer personnel during the telephone call and accepted by the caller during the telephone call. In other embodiments, the offer can be presented as an initial

decision and the offer finalised by the insurer following the call and issued in writing to the insured party to accept in writing.

On the representative display, the option to submit the information provided 120, 125, 130 (optionally along with the estimated and/or an authorised damage repair cost) to a repair shop or body shop and to authorise the repair of the vehicle 105 is also presented. In some 5 embodiments, the estimated repair cost is provided to the repair shop or body shop. In some embodiments, the estimated repair cost, or a figure based on the estimated repair cost, is provided to the repair shop or body shop as the authorised damage repair cost (i.e. repair costs in excess of the authorised amount requires further approval from the insurer before the 10 repair shop or body shop can proceed with making any repairs).

The representative can offer a total loss option if permitted (i.e. if the system determines that the vehicle is a total loss), a cash settlement option if permitted (i.e. if permitted by the insurer rules), or authorise the repair if permitted (i.e. if the vehicle is not a total loss and the claim is valid and enough information has been provided).

15 The web app 205 in this embodiment can be provided as a computer application that is executed within a standard browser and accessible using a web URL, which allows anyone with a browser capable computing device (e.g. a smartphone, tablet computer, laptop or desktop computer) to access the web app while on a call with the insurer 110. Alternatively, it can be provided as or via a native application for a smartphone, tablet computer or a 20 laptop/desktop computer.

Referring to Figure 3, an example embodiment is shown of a method of processing photos of a damaged vehicle in order to predict damage to a vehicle and estimate the repair cost for the damaged vehicle 300 and this method will now be described.

25 One or more photos of a damaged vehicle 130, typically in the region of ten to twenty photos, are provided by the caller to the insurer. These photos 130 are provided to the damage estimation process shown in this embodiment. The photos 130 will be of the damaged area of the vehicle primarily, but perhaps from a few angles and from different distances (for example, to show a close-up photo of any damage as well as a more contextual photo showing the damaged area and the surrounding undamaged portions of the vehicle).

30 In addition, vehicle information 131 is provided to the damage estimation process shown in this embodiment. Vehicle information can include model information, specifics of the vehicle (including condition, colour, optional features chosen when manufactured, modifications made to the vehicle compared to the standard). Images captures of this vehicle in this embodiment (and other embodiments) would include images showing the overall 35 vehicle (for example, images taken from at least roughly all four corners of the vehicle) as well as images showing any damage to the vehicle.

The damage estimation process uses the photos 130 and vehicle information 131 in a computer vision analysis process 305, using a trained computer vision model. The computer vision analysis process 305 identifies the parts of the vehicle shown in the photos 130, optionally using the vehicle information 131, using a set of multiple trained models, each of which trained models having been trained to recognise a specific and/or generalised vehicle part (e.g. a front bumper, or a hood, or left hand front door). The output of each of these models is combined by the computer vision analysis process 305 into a damage vector. The damage vector combines the prediction output by each trained model for each part, indicating a prediction of whether each part is damaged, and/or an indication for each part of how damaged that part is and/or in what way it is damaged, and/or along with a confidence value for each part representing the certainty of the prediction made for each part. In this embodiment, the damage vector 310 comprises the combined predictions output by the trained models 305 along with a confidence value. In other embodiments the damage vector 310 can include additional or alternative outputs from the computer vision process 305 and/or other trained models processing the input photos 130 and/or vehicle information input 131.

The damage vector 310 is output by the computer vision process 305 into (a) a representative/caller confirmation process 315 and (b) a blend analysis process 320 and (c) a repair/replace analysis process 325.

In some embodiments, where portions or all of the damage vector 310 indicate a low confidence prediction, the representative/caller confirmation process 315 displays to the representative if there are any low confidence predictions (for example, having confidence values below a predetermined threshold) for whether a part is damaged to allow the representative to guide the caller to take additional photos 130 of the one or more parts/regions of the damaged vehicle for which there are low confidence predictions. Alternatively, or in addition, the representative can be prompted during the representative/caller confirmation process 315 to ask certain questions, depending on which part has a low confidence prediction, and the representative can input the answers to the representative/caller confirmation process 315 in order to increase the prediction confidence above an acceptable threshold by augmenting the input data, allowing a revised prediction to be generated by the one or more trained models. The questions that are provided to the representative by the representative/caller confirmation process 315 can be obtained from a database (not shown) of questions that can be asked per part/region/normalised part of the damaged vehicle and/or from additional photos of the damaged vehicle.

The blend analysis process 320 includes a set of multiple trained models, each trained to perform blend analysis per part of the vehicle using the damage vector 310 and the photos 130 and the vehicle information 131, along with any further information from the representative/caller confirmation process 315. Specifically, for each part of interest the blend

analysis process 320 assesses whether any neighbouring parts are damaged and will be repainted, or replaced with a pre-painted part, in order to determine whether the part under consideration will need to have a paint blending operation performed in order to blend the paint on the part under consideration with the paint applied to the repaired or replaced part(s) neighbouring the part under consideration. In some embodiments, all parts will be considered, whereas in other embodiments only selected or determined parts will be considered by the blend analysis process.

The repair/replace analysis process 325 similarly uses a set of multiple trained models, each model trained to perform an assessment of a specific part of a vehicle (be this a specific part, or region, of normalised part of the damaged vehicle). Specifically, each model assesses, based on the input damage vector 310 and the photos 130 and the vehicle information 131, along with any further information from the representative/caller confirmation process 315, whether the part under consideration needs to be replaced or repaired.

To produce an output estimation of damage and the predicted repair cost for the vehicle, the outputs of the blend analysis process 320 and the repair/replace analysis process 325 are combined with data obtained by a part lookup process 326 by the output estimation process 330. The part lookup process 326 uses the input vehicle information 131 to determine the current market prices for each part for the vehicle under consideration and provides this to the output estimation process 330. The output estimation process 330 combines the predictions for damage to each part of the vehicle, whether each damaged part needs to be repaired or replaced, the current prices for parts to be replaced, and any painting/blending cost. The output of the output estimation process 330 is a prediction for the repair cost for the damaged vehicle. In alternative embodiments, the output is a prediction of the repairs necessary, optionally broken down into materials and/or repair operations required.

The photos 130 may be all of the photos provided to the insurer by a caller or may be a selection of those images either selected automatically (e.g. based on quality of image, image content, image metadata, a provided image description or caption from the representative or caller, or insurer rules) or manually (i.e. selected for use by the representative or other personnel of the insurer, or highlighted by the representative or user as relevant).

The vehicle information 130 may sometimes be only a very limited amount of information such as the make and model of the vehicle, its age and government registration details, and its paint colour, or it can be a richer dataset including comprehensive details about the vehicle specification and condition as well as caller-provided and insurer-collected information about the vehicle. Different embodiments have different minimum input data requirements in respect of the vehicle information 130.

Different trained models and trained model arrangements can be used in the computer vision analysis process 305 and/or the blend analysis 320 and or the repair/replace analysis 325, and different arrangements for these models can be implemented such as to train one or a plurality of models that can recognise two or more parts rather than just one specific part as per the above described embodiment.

During the representative/caller confirmation process 315, the information on low confidence predictions from the damage vector 310 can be presented to either the representative or directly to the caller via the web app in order to obtain either further photos from the caller or information to increase the confidence of the predictions for which there is low confidence output from the computer vision process 305.

In alternative embodiments, there is optionally an additional paint estimation process included as part of the method of processing photos of a damaged vehicle in order to predict damage to a vehicle and estimate the repair cost for the damaged vehicle. The inputs can include some of the outputs from other processes. The paint estimation process can use multiple trained models to assess, per part, whether any painting is required for each part of the damaged vehicle or can obtain this data from a third-party database. The output of these models is provided to the output estimation process 330 which uses this additional information to output the prediction for the repair cost for the damaged vehicle.

In alternative embodiments, there is optionally an additional labour estimation process as part of the method of processing photos of a damaged vehicle in order to predict damage to a vehicle and estimate the repair cost for the damaged vehicle. The labour estimation process can use multiple trained models to determine, per part, the extend of the damage and whether any labour is required to perform repairs and/or replace any parts. A lookup can then be performed to produce an output value by using the determination to obtain this output data from a third-party database, for example to determine an output value by submitting a query for a time value for one or more labour tasks to repair and/or replace parts along with some or all of the vehicle information input 131 and receiving in return from the third-party database repair information determined for that job for the specific make, model and year (MMY) of the damaged vehicle. Some of the outputs from other estimation processes can be used as inputs to the labour estimation process. The output of these models is provided to the output estimation process 330 which uses this additional information to output the prediction for the repair cost for the damaged vehicle.

In alternative embodiments, there is optionally an additional paint labour estimation process as part of the method of processing photos of a damaged vehicle in order to predict damage to a vehicle and estimate the repair cost for the damaged vehicle. The paint labour estimation process can use multiple trained models to determine, per part, whether no paint is needed, a "spot" paint job is needed, or more or less than 50% of the part needs to be

5 painted. A lookup can then be performed to produce an output value by using the determination to obtain this output data from a third-party database, for example to determine an output value by submitting a query for a "spot" paint job along with some or all of the vehicle information input 131 and receiving in return from the third-party database repair information determined for that job for the specific make, model and year (MMY) of the damaged vehicle. Some of the outputs from other estimation processes can be used as inputs to the labour estimation process. The output of these models is provided to the output estimation process 330 which uses this additional information to output the prediction for the repair cost for the damaged vehicle.

10 In some embodiments, there is optionally an additional remove and installation estimation process as part of the process of estimating a repair cost for the damaged vehicle. This estimation process estimates the amount of time that is necessary to perform all of the tasks required for replacing a damaged part of a vehicle. These estimates are obtained from a rule set for each part or group of parts/type of part that has been built up by an expert based on knowledge of repairs to vehicles, or are obtained from a third party data provider to which a query specifying the part in question and the vehicle information and some damage information is transmitted and an estimate returned by the data provider. Optionally, this process receives one or more outputs from other modules that are included in its output. The output from the rule set or the returned estimate from the data provider is then used as the output from the remove and installation process.

15 In other embodiments, there is optionally an additional overlap calculation performed. The overlap that can be identified is the overlap in repair work to be performed when considering repairs to neighbouring parts that are damaged, as repairs to multiple neighbouring parts will involve common components that only need to be replaced/handled or tasks that can be performed once and thus do not need to be repeated for each part. In some embodiment, this overlap is determined from a query to a third-party data provider. Thus, a modifier can be determined based on the number and degree of overlap between repair tasks and components to take into account the overlap between repairs to neighbouring damaged parts of a vehicle.

20 Referring to Figure 4, an example embodiment 400 of a method of predicting total loss based on the input information and photos provided by a caller to the insurer and the representative will now be described.

25 In this embodiment, a method 400 of predicting total loss is described that can be used with any of the other described embodiments or alternative embodiments, or with the aforementioned aspects described in this specification.

According to this embodiment, the method 400 takes as input data the crash and insurance details 405 and the repair estimate 410 output by the damage prediction model of one of the embodiments described herein and the vehicle value 415.

5 The crash and insurance details 405 will be provided by the caller to the insurer's representative when calling or contacting the insurer to report the accident/damage to the damaged vehicle as described in relation to the other described embodiments herein.

10 The estimate for repair costs for the damaged vehicle 410 is determined using the damage repair estimation process described in the embodiments and alternative embodiments that have been described herein. This estimate 410 may comprise any combination of: the operations required to repair the damaged vehicle; the parts and/or materials required to repair the damage vehicle; and the associated financial costs of repairing the damaged vehicle.

15 The vehicle value 415 may be estimated, and/or have been provided by the insured party, and/or be determined by the insurer using rules or other guidelines, and/or is provided by a third party.

20 The prediction of total loss 420 is then determined, using a trained model, from a variety of factors including whether the crash details 405 indicate a class of accident that typically results in significant structural damage to a car making it unsafe. In other embodiment, other factors/combinations of factors can be used to make this determination. If it is determined that there has been significant structural damage, for example, the assumption can be made that the accident is a total loss. In alternative embodiments, the details of the accident can be assessed using a database of classes of accident (the database can be generated using knowledge from experts, or using data from insurers, or accessed through the use of a third part data provider) and/or the damage assessment that forms part of the repair estimate 410 to determine whether there is a high likelihood of a total loss or whether there is likely to be significant non-visible damage to the vehicle. In other embodiments, the repair estimate is generated (including an estimate of the financial costs of the repair) as per the described embodiments and if the total value of the repair is above a certain value, for example the replacement value or insured value of the vehicle, or a substantial fraction thereof, then it is determined that the accident has resulted in a total loss. The determination of total loss 420 also considers the damage repair estimate 410 against the vehicle value 415 using the insurer's own rules, policies and/or procedures to determine whether the cost of the repairs 410 is likely to exceed the vehicle value 415 or a significant proportion of the vehicle value 415.

35 Referring to Figure 5, another example embodiment is shown of a method 500 of processing photos of a damaged vehicle in order to predict damage to a vehicle and estimate the repair cost for the damaged vehicle, and this method 500 will now be described.

One or more photos of a damaged vehicle 130, typically in the region of ten to twenty photos, are provided by the caller to the insurer representative. These are provided to the damage estimation process 500 shown in this embodiment. The photos 130 will be of the damaged area of the vehicle primarily, but typically from a few different viewing angles and from different distances (for example, to show a close-up photo of any damage as well as a more contextual photo showing the damaged area and the surrounding undamaged portions of the vehicle).

In addition, vehicle information 131 is provided to the damage estimation process shown in this embodiment. Vehicle information 131 can include model information, specifics of the vehicle (including condition, colour, optional features chosen when manufactured, modifications made to the vehicle compared to the standard).

The damage estimation process uses the photos 130 and the vehicle information 131 in a computer vision analysis process 305. The computer vision 305 identifies the parts of the vehicle shown in the photos 130, optionally using the vehicle information 131, and uses a set of multiple trained models, each of which trained models has been trained to recognise a specific generic vehicle part or region (e.g. a front bumper, or a hood, or left hand front door). The output of each of these per part/region models is combined by the computer vision analysis process 305 into a damage vector 310. The damage vector 310 combines the prediction output by each trained model for each part/region indicating a prediction of whether each part/region is damaged, and an indication for each part/region how damaged that part/region is, along with a confidence value for each part/region representing the certainty of the prediction made for each part/region. In other embodiments, specific parts for a specific make, model and year (MMY) of a vehicle are recognised instead of regions/generalised parts/normalised parts/generic regions, optionally by using both the photos 130 and the vehicle information 131 rather than just the photos 130.

The damage vector 310 is output by the computer vision process 305 into (a) a representative/caller confirmation process 315 and to (b) a blend analysis process 320 and (c) a repair/replace analysis process 325 and (d) a paint analysis process 505. Alternatively, the paint analysis process 505 can have its input as one or more of the outputs from one or more of the other processes 320, 325.

The representative/caller confirmation process 315 displays to the representative if there are any low confidence predictions (i.e. under a predefined threshold confidence value) for whether a part/region is damaged, to allow the representative to guide the caller to take additional photos 130 of the one or more parts/regions for which there are low confidence predictions. Alternatively, or in addition, the representative can be prompted during the representative/caller confirmation process 315 to ask certain questions, depending on which part/region has a low confidence prediction, and the representative can input the answers to

the representative/caller confirmation process 315, which information is provided to the model 305 in order for the model to recalculate any low-confidence predictions, in order to increase the prediction confidence above an acceptable threshold. The questions that are provided to the representative by the representative/caller confirmation process 315 can be obtained from a database (not shown) of questions that can be asked per part/region of interest.

The blend analysis process 320 includes a set of multiple trained models, each trained to perform blend analysis per part/region of the vehicle using the damage vector 310 and the photos 130, along with any further information obtained from the representative/caller confirmation process 315. Specifically, the blend analysis process 320 assesses whether any neighbouring parts/regions are damaged, and thus will be repainted or replaced with a pre-painted part, in order to determine whether the part/region under consideration will need to have a paint blending operation performed in order to blend the paint on the part/region under consideration with the paint applied to the repaired or replaced part/regions(s) neighbouring the part/region under consideration.

The paint estimation process 505 uses multiple trained models to assess, per part/region, whether any painting is required for each part/region of the damaged vehicle. The input to the paint estimation process 505 is the damage vector 310 (and, in some embodiments, the photos 130), along with any further information obtained from the representative/caller confirmation process 315. The output of these models is provided to the output estimation process 330.

The repair/replace analysis process 325 again includes a set of multiple trained models, each trained to perform an assessment of a specific part/region of a vehicle. Specifically, each model assesses, based on the input damage vector 310 and the photos 130, along with any further information obtained from the representative/caller confirmation process 315, whether the part/region under consideration needs to be replaced or repaired.

To produce an output estimation of damage and the predicted repair cost for the vehicle, the outputs of the blend analysis process 320 and the repair/replace analysis process 325 and the paint analysis process 505 are combined with data obtained by a part lookup process 326 by the output estimation process 330. The part lookup process 326 uses the input vehicle information 131 to determine the current market prices for each part, or associated parts for the region, for the vehicle under consideration and provides this to the output estimation process 330. The output estimation process 330 combines the predictions for damage to each part/region of the vehicle, whether each damaged part/region needs to be repaired or replaced, the current prices for parts to be replaced, and any painting/blending cost. The output of the output estimation process 330 is the prediction for the repair cost for the damaged vehicle.

The photos 130 may be all of the photos provided to the insurer by a caller or may be a selection of some of those images, either selected automatically (e.g. based on quality of image, image content, image metadata, a provided image description or caption from the representative or caller, or insurer rules) or manually (i.e. selected for use by the representative or other personnel of the insurer, or highlighted by the representative or user as relevant).

The vehicle information 130 can be only a very limited amount of information such as the make and model of the vehicle, its age and government registration details, and its paint colour, or it can be a richer dataset including comprehensive details about the vehicle specification and condition as well as caller-provided and insurer-collected information about the vehicle.

Different trained models and trained model arrangements can be used in the computer vision analysis process 305, and different arrangements for these models can be implemented such as to train one or a plurality of models that can recognise a generalised part/region, or two or more parts, rather than just one specific part in the above described embodiments.

During the representative/caller confirmation process, the information on low confidence predictions from the damage vector 310 can be presented to either the representative or directly to the caller via the web app in order to obtain either further photos from the caller or information to increase the confidence of the predictions for which there is low confidence output from the computer vision process 305.

Alternatively, a further labour estimation process (not shown) can optionally also be included in embodiments and used as part of the output estimation process 330 to determine a prediction for the repair cost for the damaged vehicle.

In some embodiments, there is also a remove and installation estimation process (not shown) optionally included as part of the process of estimating a repair cost for the damaged vehicle. This estimation process estimates the amount of time that is necessary to perform all of the tasks required for replacing a damaged part of a vehicle. These estimates are obtained from a rule set for each part or group of parts/type of part that has been built up by an expert based on knowledge of repairs to vehicles, or are obtained from a third party data provider to which a query specifying the part in question and the vehicle information and some damage information is transmitted and an estimate returned by the data provider. The output from the rule set or the returned estimate from the data provider is then used as the output from the remove and installation process.

In other embodiments, there is optionally also an overlap calculation performed (not shown). The overlap that can be identified is the overlap in repair work to be performed when considering repairs to neighbouring parts that are damaged, as repairs to multiple neighbouring parts will involve common components that only need to be replaced/handled or

tasks that can be performed once and do not need to be repeated for each part. Thus, a modifier can be determined based on the number of, and degree of overlap between, repair tasks and components to take into account the overlap between repairs to neighbouring damaged parts/regions of a vehicle.

5 Referring to Figure 6, an example embodiment of a process 600 of taking the photos 130 provided by the caller and processing these to output a damage signature vector 535 will now be described.

As in the other described embodiments, the caller provides photos 130 of the damaged vehicle. Typically these might number in the region of 10 to 20 images, but fewer photographs
10 may be supplied (for example where only minor damage is visible) or conversely more photographs may be supplied (for example where there is extensive damage that requires many photographs to document and/or the caller provides a comprehensive selection of close up and distant/further removed photos of the damaged vehicle).

The part classifier 505 is a set of multiple models 505i, 505ii, 505iii, 505iv to 505n (n
15 being the number of models) each of which runs in parallel on each of the photos 130 and each of which models 505i, 505ii, 505iii, 505iv to 505n is trained to determine whether a photograph contains a specific part/region of a vehicle. The output of each of the models 505i, 505ii, 505iii, 505iv to 505n are the determinations whether each part/region (for which each model is trained to detect) is present in each of the photos and some metadata, for example
20 optionally indicating the damage predicted for that part and/or an associated confidence value.

The output from each of the multiple models 505i, 505ii, 505iii, 505iv to 505n of the part classifier 505 is fed into each of a set of further assessment models per part along with the photos 130. To illustrate the next stages in the process/method, only the output from one model 505i will be described here but the reader will understand that each of the multiple
25 models 505i, 505ii, 505iii, 505iv to 505n of the part classifier 505 creates an output that is fed to a like group of assessment models per part in parallel with the described process in relation to the output from one model 505i that will now be described.

The output from the first models 505i of the part classifier 505 and the photos 130 is fed into multiple further assessment models, specifically a paint assessment model 511, a
30 blend assessment model 510, a repair or replace determination model 515 and a labour estimation model 520.

The paint assessment model 511 receives the photos 130 and determines which of the photos 130 are deemed relevant by the part classifier model 505i (i.e. show the part of interest) along with the output metadata indicating the damage predicted for that part and an
35 associated confidence value. The paint assessment model 511, for each relevant photo of the input photos 130, then predicts whether and the extent to which painting is required for each photo of the damaged part. This prediction is then output for each relevant photo of the

damaged part and an initial estimate is determined 513 based on a pooled score of all of the predictions output by the paint assessment model 511 for all of the photos of the damaged part provided as an input to the paint assessment model 511 by the part classifier model 505i. Alternatively, the paint assessment model 511 can receive the outputs from other models and the paint assessment process submits queries to one or more third party databases to obtain information on what paint costs and tasks are required.

The blend assessment model 510 receives the photos 130 and determines which of the photos 130 are deemed relevant by the part classifier model 505i along with the output metadata indicating any damage predicted for that part and neighbouring parts and an associated confidence value. The blend assessment model 510, for each relevant photo of the input photos 130, then predicts whether and the extent to which blending is required for each photo of the part being considered based on the damage predicted for that or any neighbouring parts. This prediction is then output for each relevant photo of the part being considered and an initial estimate is determined 512 based on a pooled score of all of the predictions output by the blend assessment model 510 for all of photos of the part being considered provided as an input to the blend assessment model 510 by the part classifier model 505i.

The repair or replace determination model 515 receives the photos 130 and determines which of the photos 130 are deemed relevant by the part classifier model 505i along with the output metadata indicating the damage predicted for that part and an associated confidence value. The repair or replace determination model 515, for each relevant photo of the input photos 130, then predicts whether repair is possible or whether replacement is necessary for each photo of the damaged part based on the damage predicted by the part classifier model 505i for that photo. This prediction is then output for each relevant photo of the damaged part and an initial estimate is determined 517 based on a pooled score of all of the predictions output by the repair or replace determination model 515 for all of the photos of the damaged part provided as an input to the repair or replace determination model 515 by the part classifier model 505i.

The labour estimation model 520 receives the photos 130 and determines which of the photos 130 are deemed relevant by the part classifier model 505i along with the output metadata indicating the damage predicted for that part and an associated confidence value. The labour estimation model 520, for each relevant photo of the input photos 130, then predicts what labour for repair or replacement or painting of the part under consideration for each photo of the part under consideration based on the damage predicted by the part classifier model 505i for that photo and any neighbouring parts. This prediction is then output for each relevant photo of the part under consideration and an initial estimate is determined 522 based on a pooled score of all of the predictions output by the labour estimation model

520 for all of the photos of the part under consideration provided as an input to labour estimation model 520 by the part classifier model 505i. Alternatively, based on the output of one or more of the other processes, the labour estimation model 520 submits queries to one or more third party databases to obtain labour task and price estimates.

5 The output initial estimates 513, 512, 517, 522 are then concatenated 530, along with any other initial estimates produced for painting, blending, repair, replacement or labour for each of the other parts processed by the multiple models 505i, 505ii, 505iii, 505iv to 505n of the part classifier 505. The output of the concatenation 530 is a vector with classification 535 (otherwise known as a damage signature).

10 Referring to Figure 7, an example embodiment of a blending assessment model 700 will now be described.

The inputs to the blending assessment model 715 are the damage signature output from the computer vision or part detection model 535 along with the car details 705 and the information about the damage to the vehicle 710 provided by the caller to the insurer representative.

15 The blending assessment model 715 works on a per part basis, so is trained to assess whether blending is required for a specific part/region of a vehicle and uses information in the damage signature 535 on what damage there is to neighbouring parts to assess whether an undamaged part/region requires blending with the new paint applied to either a repaired part or a replaced part that is painted or comes pre-painted. The model 715 is a trained machine learned model in this embodiment.

20 The output 720 of the blending assessment model 715 is a binary output per part as to whether blending is required for each part that is considered by each of the multiple per part blending assessment models 715. In other embodiments, more detailed outputs can be produced.

25 Referring to Figure 8, an example embodiment 800 of a repair or replace model (per part) 715 will now be described.

30 The repair or replace model (per part) 715 works on a per part basis, so is trained to assess whether a repair to each part is possible or whether a replacement part is required for a specific part of a vehicle. The repair or replace model (per part) 715 uses information in the damage signature 535 on what damage there is to the part being considered to assess whether a part is undamaged or whether a part can be repaired or requires replacing with a new part. The model 715 in this embodiment is a trained machine learned model.

35 The output 720 of the repair or replace model (per part) 715 is an output per part indicating whether each part is considered by each of the multiple per part estimation models 715 to be undamaged, repairable or need replacing.

Referring to Figure 9, an example embodiment 900 of a variant of the method of producing a repair estimate from photos 130 of a damaged vehicle will now be described.

One or more photos of a damaged vehicle 130, typically in the region of ten to twenty photos, are provided by the caller to the insurer representative. These are provided to the damage estimation process shown in this embodiment. The photos 130 will be of the damaged area of the vehicle primarily, but perhaps from a few viewing angles and from different distances (for example, to show a close-up photo of any damage as well as a more contextual photo showing the damaged area and the surrounding undamaged portions of the vehicle).

The damage estimation process first uses the photos 130 in a computer vision analysis process 305. The computer vision analysis process 305 identifies the parts/regions of the vehicle shown in the photos 130 using of a set of multiple trained models, each of which trained models has been trained to recognise a specific vehicle part/region (e.g. a front bumper, or a hood, or left hand front door). The output of each of these models is combined by the computer vision analysis process 305 into a damage vector 310. The damage vector 310 combines the prediction output by each trained model for each part/region indicating a prediction of whether each part/region is damaged, and an indication for each part/region how damaged that part/region is, along with a confidence value for each part/region representing the certainty of the prediction made for each part/region.

The damage vector 310 is output by the computer vision process 305 into (a) a representative/caller confirmation process 315 and to (b) a blend engine 320 and (c) a repair/replace analysis process 325 and (d) a remove and install (R&I) process 910. In some embodiments, the damage vector 310 is also provided to a paint analysis process 505 and to a labour analysis process. In this embodiment, the outputs of the other processes 320, 910, 325 are provided as inputs to the repair labour analysis process 905 and paint analysis process 505.

The representative/caller confirmation process 315 displays to the representative if there are any low confidence predictions within the damage vector 310, as to whether a part/region is predicted to be damaged or not, to allow the representative to guide the caller to take additional photos 130 of the one or more parts for which there are low confidence predictions. Alternatively, or in addition, the representative can be prompted by the representative/caller confirmation process 315 to ask certain questions, depending on which part has a low confidence prediction, and the representative can input the answers to the representative/caller confirmation process 315 in order to increase the prediction confidence above an acceptable threshold. The questions that are provided to the representative by the representative/caller confirmation process 315 can be obtained from a database (not shown) of questions that can be asked per part. The output from this process is provided to the representative/caller confirmation process 315 and the revised output of the computer vision

process 305, once the new information obtained during the confirmation process 315 has also been obtained, is a revised damage vector 310 which is provided to the blend engine 320, the repair/replace analysis process 325, the paint analysis process 505, and to the labour analysis process (via the blend engine 320). Alternatively, the output from this process is provided
5 along with the original damage vector 310 to the representative/caller confirmation process 315, the blend engine 320, the repair/replace analysis process 325, the paint analysis process 505, and to the labour analysis process (via the blend engine 320).

The blend engine 320 (or blend analysis process) includes a set of multiple trained models, each trained to perform blend analysis per part/region of the vehicle using the damage
10 vector 310 and the photos 130, along with any further information from the representative/caller confirmation process 315. Specifically, the blend analysis process 320 assesses whether any neighbouring parts/regions are damaged and will be repainted or replaced with a pre-painted part in order to determine whether the part/region under consideration will need to have a paint blending operation performed in order to blend the
15 paint on the part under consideration with the paint applied to the repaired or replaced part/region(s) neighbouring the part/region under consideration.

The paint estimation process 505 uses multiple trained models to assess, per part/region, whether any painting is required for each part/region of the damaged vehicle. In some embodiments, the input to the paint estimation process 505 can be the damage vector
20 310 and the photos 130, along with any further information from the representative/caller confirmation process 315 and/or a revised damage vector 310. In this embodiment, the input is the outputs per part from the blend engine 320, the R&I analysis 910 and the repair/replace analysis 325. The output of these models is provided to the output estimation process 330. The paint estimation process 505 determines whether, for each part, no paint is needed, a
25 "spot" paint job is needed, or more or less than 50% of the part needs to be painted. To obtain an estimate of the hours of labour needed to perform the determined level of painting, and to obtain a price for the paint required, one or more queries are sent to one or more third party databases to obtain data or an estimate.

The repair/replace analysis process 325 again includes a set of multiple trained
30 models, each trained to perform an assessment of a specific part/region of a vehicle. Specifically, each model assesses, based on the input damage vector 310 and the photos 130, along with any further information from the representative/caller confirmation process 315, whether the part/region under consideration needs to be replaced or repaired.

The repair labour analysis module 905 takes input from the blend engine 320 and the
35 repair/replace assessment module 325 and the paint analysis module 505 and makes a determination of the required labour time/amount needed to make the determined/estimated

repairs/replacements and any paint application, and this determination is output to the estimation output process 305 for incorporation into the predicted repair cost.

In some embodiments, there is also an optional remove and installation estimation process 910 as part of the process of estimating a repair cost for the damaged vehicle. This estimation process 910 estimates the amount of time that is necessary to perform all of the tasks required for replacing a damaged part/region of a vehicle. These estimates are obtained from a rule set for each part or group of parts/type of part that has been built up by an expert based on knowledge of repairs to vehicles, or are obtained from a third party data provider to which a query specifying the part in question and the vehicle information and some damage information is transmitted and an estimate returned by the data provider. The output from the rule set or the returned estimate from the data provider is then used as the output from the remove and installation process.

In other embodiments, there is also an optional overlap calculation performed. The overlap that can be identified is the overlap in repair work to be performed when considering repairs to neighbouring parts that are damaged, as repairs to multiple neighbouring parts will involve common components that only need to be replaced/handled or tasks that can be performed once and do not need to be repeated for each part. Thus, a modifier can be determined based on the number and degree of overlap between repair tasks and components to take into account the overlap between repairs to neighbouring damaged parts of a vehicle. In this embodiment, calculation of any modifications to values due to any overlap is performed by the repair labour analysis module 905.

To produce an output estimation of damage and the predicted repair cost for the vehicle, the outputs of the blend analysis process 320 and the repair/replace analysis process 325 and the paint analysis process 505 and the repair labour analysis module 905 are combined by the output estimation process 330 and any modified calculation is applied. The output estimation process 330 combines the predictions for damage to each part of the vehicle, whether each damaged part needs to be repaired or replaced, any painting/blending cost and the labour cost (all of some of which having had a modifier, if calculated, applied to reflect a reduction due to overlap). The output of the output estimation process 330 is the prediction for the repair cost for the damaged vehicle.

In some embodiments, a overlap engine is used to perform an overlap calculation, to determine one or more reductions to be applied, for example where two neighbouring parts require repair operations to be performed, a reduction in labour time due to efficiencies from performing the repair operations/labour at the same time for both neighbouring parts can be applied to the overall total repair cost/operations/materials.

The photos 130 may be all of the photos provided to the insurer by a caller or may be a selection of those images either selected automatically (e.g. based on quality of image,

image content, image metadata, a provided image description or caption from the representative or caller, or insurer rules) or manually (i.e. selected for use by the representative or other personnel of the insurer, or highlighted by the representative or user as relevant).

5 Different trained models and trained model arrangements can be used in the computer vision analysis process 305, and different arrangements for these models can be implemented such as to train one or a plurality of models that can recognise two or more parts rather than just one specific part as per the above described embodiment.

10 During the representative/caller confirmation process, the information on low confidence predictions from the damage vector 310 can be presented to either the representative or directly to the caller via the web app in order to obtain either further photos from the caller or information to increase the confidence of the predictions for which there is low confidence output from the computer vision process 305.

15 In some embodiments, vehicle information 131 can also be provided to the damage estimation process shown in this embodiment. Vehicle information can include model information, specifics of the vehicle (including condition, colour, optional features chosen when manufactured, modifications made to the vehicle compared to the standard). The vehicle information 130 can be only a very limited amount of information such as the make and model of the vehicle, its age and government registration details, and its paint colour, or it can be a
20 richer dataset including comprehensive details about the vehicle specification and condition as well as caller-provided and insurer-collected information about the vehicle. The vehicle information can be used by any of the computer vision module 305, the blend engine 320, the paint estimation module 505, the repair/replace analysis module 325 and the repair labour analysis module 905 as well as for producing the final estimation output 330. A part lookup
25 process 326 can use the input vehicle information 131 to determine the current market prices for each part for the vehicle under consideration and provide this to the output estimation process 330.

Referring to Figure 10, an example embodiment 1000 of a trained model architecture 1030 will now be described.

30 The trained model architecture 1030 is a neural network with a series of layers 1060, 1070, 1080, 1090, 1100 between most of which layers there are injection points 1075, 1085, 1095, 1105.

The input to the neural network 1030 is an image of a part 1010. The neural network 1030 is also provided with metadata 1020, which is split into two groups of metadata 1040,
35 1050, and which metadata 1040, 1050 is injected into the neural network 1030 between layers 1070, 1080, 1090, 1100 at injection points 1075, 1085, 1095, 1105.

The metadata 1020, 1040, 1050 can include (i) the repair/replace scores predicted for the part/region in question; (ii) the repair/replace scores predicted for any, multiple or all nearby parts/regions (iii) the undamaged/damaged values of the part/region in question (iv) the undamaged/damaged values of any, multiple or all nearby parts/regions; (v) the type of vehicle (e.g. whether the vehicle is a pick-up truck, car, van, etc); (vi) the number of doors the vehicle has (e.g. whether the vehicle has two, three or four doors); and (vii) the colour of the vehicle (e.g. whether the vehicle is white, or where there are specific colours to types of vehicle, more precise colours or colour metadata can be captured). Other metadata can be injected in other embodiments, for example details about the vehicle (make, model and year; options selected, etc), damage (for example, the output of any classifiers/models), and accident circumstances (for example, prior damage, circumstances of the cause of the damage, etc).

In an embodiment where the neural network 1030 is being used as a blending assessment model per part/region, only items (ii), (iv), (v) (vi) and (vii) of metadata 1020, 1040, 1050 will be used.

Each of items (ii), (iv), (v) (vi) and (vii) of metadata 1020, 1040, 1050 are injected separately at each of the injection points 1075, 1085, 1095, 1105.

The output 1200 of the neural network 1030 depends on how it has been configured. In an embodiment where the neural network 1030 is configured as a blending assessment model, the output will be a binary yes or no answer as to whether the part in question requires a paint blending operation to be performed thereupon.

Alternatively, instead of the neural network being constructed from a series of single layers 1060, 1070, 1080, 1090, 1100, the neural network can be constructed from one or more groups of layers.

In other embodiments the neural network 1030 can be configured to be a blend engine, a repair/replace analysis process, a remove and install process, a modifier calculation process, a paint analysis process, or a labour analysis process. Different metadata 1020 will be used in each of these embodiments.

As shown in Figure 11, which will be described in more detail below, conventionally vehicle repair cost estimates (or “claim estimates”) are generated at a body shop (i.e. a garage or workshop where vehicles are serviced and/or repaired and which has been proposed by either a customer or insurer as the repairer of the insured damaged vehicle) when a customer brings in a damaged vehicle that is insured. The estimates are then reviewed by the insurer (this is typically known as the “claim review process”). Insurers, specifically a specialist team of people or representatives working for the insurer, when reviewing claim estimates, may be tasked to follow pre-defined procedures or workflows. Such procedures or workflows may require the relevant people/representatives to routinely challenge claim estimates where

statistical anomalies are identified, for example. However, insurers face issues in reviewing claims accurately and efficiently as the task when performed manually is time-intensive and prone to human error. Inefficient processes in reviewing claims can result in a backlog of claims to be reviewed manually and can accumulate high costs associated with the inefficient processes. Providing an automated or semi-automated claim review system can help triage claims, and can help to identify leakage (i.e. proposed repairs that are either not covered by the relevant insurance policy or are unnecessary, incorrectly priced or even fraudulent) in claims, using automated assessments of the data and information provided to the insurer by both the insured party and the proposed repairer.

With reference to Figures 12 to 17, example embodiments relating to assessing claim input data using an automated assessment platform for verifying the claim input data, by generating estimates of damages and determining whether claim input data can be approved or is anomalous or whether further assessment may be required, will now be described.

Example embodiments provide an automated system to assess or review whether vehicle repair estimates are acceptable and legitimate. The system can be used to flag repair operations that are likely to be irregular in some way (i.e. potential leakage), which can then be further reviewed by the insurer (usually manually, by a specialist claims reviewer or loss adjuster).

In some embodiments, where no irregular operations are identified, or the number (or value) of irregular operations are below a pre-determined threshold, the repair operations can be automatically approved on behalf of the insurer, or allocated to a queue or workflow of the insurer.

As shown in Figure 11, which shows a traditional vehicle repair claim process 1100, in cases of vehicle damage or following accidents 1105, the insured party will typically either take the insured vehicle to a repair shop or body shop for repair 1115 and will contact their insurer to report the accident 1110.

The repair shop or body shop may be chosen by the insured party or by the insurer, depending on various factors such as but not limited to the terms of the relevant insurance policy and the preference of the insured party.

If the insured party takes the damaged vehicle to the repair shop 1115 first, the damage is assessed by the repair shop in order for them to prepare a quote for, or an estimated cost of, the repair 1120 that is required to the damaged vehicle. Depending on the various factors, such as the procedures of the repair shop, the policy of the insurer and the severity of the damage or accident, the estimate from the repair shop is sent to the insurer 1125 for authorisation by the insurer as typically insurers scrutinise any repair estimates received 1130. The repair shop may transmit their estimated repair cost along with supporting details (e.g. repair labour proposed, parts required, materials required such as paint) to the insurer 1125.

In addition to the repair estimate 1120 being sent to the insurer by the repair shop, the party that is insured (or a party reporting the damage/accident on behalf of an insured entity or in respect of an insured vehicle) will contact their insurer 1110 to report the accident or damage that has occurred and provide details to the insurer about the accident or damage 1125 including providing images of the damaged vehicle (which are provided either by the driver or the repair shop).

The insurer will then review the claim, optionally requesting further details from the insured party and/or the repair shop as necessary to verify the claimed amounts before approving the repair work or sending payment to either the insured party or the repair shop.

Specifically, the representative taking the telephone call on behalf of the insurer, when the client reports the accident 1105 to the insurer 1110, may run through a standard set of questions, or a decision tree of questions, to obtain a standard required set of information about the accident (e.g. the name of the insured person/vehicle; details of the vehicle including make, model and condition; details of any insurance policy; standard security questions; details of what happened to cause the accident; etc). In this way the insurer can be provided with information, for example a standard checklist of information requested by the insurer representative by following script, that will ensure that comprehensive details of the insured party and vehicle will be provided to the insurer as well as comprehensive accident and damage details. Additionally, photographs of the damaged vehicle may also be provided to the insurer.

As the details of the insured person, accident details and photographic images of any damage are provided by the party to the insurer, the insurer can consider 1130 and scrutinise the details provided typically to assess, but not limited to, the following (a) whether there is a valid insurance policy in place; (b) whether the reported damage is likely to have occurred in a way and timeframe covered by the insurance policy; (c) whether the proposed repairs are acceptable under the insurance policy or the criteria of the insurer; (d) whether the part prices quoted are within an acceptable threshold pricing for each part; (e) whether the proposed labour costs and time for repairs and/or painting are appropriate for the determined damage to the vehicle; (f) whether each decision to repair or replace a part was correct; (g) whether the paint work on parts should be blended; (h) whether the part is an OEM or non-OEM part (i) whether a pre-painted part should or shouldn't be used; and (j) whether the correct option for a part has been specified. Sometimes, the report to the insurer 1110 or the repair estimate provided to the insurer 1125 does not include sufficient information or detail when it is considered by the insurer 1130 and so, if the insurer determines that insufficient information or detail has been provided 1135, the insurer requests further information 1140. The insurer can then reconsider the data provided 1130 again, make a similar determination 1135 until sufficient detail is provided and then assess the repair estimate 1145 and then reject the

estimate 1150 or approve the estimate 1155 so that the repair work can be performed 1160 by the garage. If the estimate is rejected 1150, then the process is repeated when the garage re-submits a revised repair estimate from step 1120.

Referring to Figure 12, an example embodiment of a claim process 1200 with an incorporated automated repair cost estimate review process will now be described.

Following an accident or damage to a vehicle 1105, the insured party/client reports the accident to their insurer 1110 following the accident 1105 and takes the vehicle to the body shop for a repair assessment 1115. The body shop prepares a repair estimate 1120 for the damage to the vehicle and reports this estimate to the insurer 1125.

The data or information in relation to a particular claim is input into an automated review system 1205, either manually by personnel at the insurer entering it into the front end of the review system or automatically by passing data received from the client and repair shop directly into the automated review system 1205.

In some embodiments the automated review system 1205 can carry out different or a combination of tasks depending on the type of data input into the automated system 1205. For example, in the embodiment shown in Figure 12, the automated review system 1205 triggers one of three review processes: a pre-defined workflow to be completed by a user 1210, the automatic approval of claims with no determined anomalies 1220; or the flagging of cases where the automated system 1205 does not agree with the repair estimate produced by the body shop 1230 so triggers a manual (i.e. human expert) or further review of such flagged cases 1235 – and the expert either agrees and the estimate is finalised 1225 or the expert disagrees and the garage is required to prepare a revised estimate 1120. Thus, in this example embodiment, the automated review system 1205 can approve, query or reject the body shop repair estimate in order to finalise the repair estimate 1225 for approval.

In the example embodiment, of the automated review system 1205 is able to perform a review by determining an independent repair estimate for the claim being reviewed based on the claim inputs and using estimatics providers to determine an independent repair estimate using a computer-generated damage estimate (generated from the claim inputs) and comparing this independent repair estimate to the provided repair estimate from the garage 1125.

In example embodiments, workflows such as checking the VIN, checking if required photos are missing, and checking that invoices match repair estimates can be performed. For example, one workflow can be triggered when it is determined that there is an error or inconsistency with the VIN provided for the vehicle, or if the VIN hasn't been provided for the vehicle – the workflow can, instead of requiring the claim to be reviewed manually, automatically return the repair estimate to the garage and request the VIN details for the vehicle before the repair estimate is considered for approval. Another example workflow would

be triggered should it be determined that insufficient photos have been provided of the damage to the vehicle, and requesting further or better quality photos of the vehicle or specific areas of damage to the vehicle from the garage before the repair estimate is considered for approval. A further example workflow would be triggered should an invoice be submitted for repair costs that do not match what was approved in a previous repair estimate, either triggering a manual (human) review or sending an automatic rejection of the invoice to the garage providing details of the inconsistency with a previously submitted repair estimate.

In embodiments, the damage to the vehicle is determined from the photos of the damaged vehicle to determine the extent of the damage and the severity of each piece of damage to the vehicle so it is important to have sufficient image data of the vehicle and that this image data is of sufficient quality. Some damage to a vehicle, such as structural damage cannot be determined from photos of the damaged vehicle, however.

With reference to Figure 13, an overview of process 1300 performed by the automated review system 1205 shown in Figure 12 will now be described in further detail.

The inputs 1305 provided to the automated assessment platform 1310 include the information provided to the insurer by the client 1110 and the repair estimate 1125 from the garage. In the example embodiment, the automated assessment platform 1210 is provided with photographs 1320 relevant to the claim (i.e. photographs of the vehicle, including for example photographs of the vehicle from a few different angles around the vehicle as well as close up photographs, generally of the damaged portions and parts of the vehicle), an order sheet 1325 prepared by the insurer (i.e. details of the accident and/or damage prepared by the insurer having spoken to the client) and the financial estimate provided by the repair shop 1330 (i.e. a list of the parts needing replacement, the materials required such as paint and a list of the labour required to remove and replace parts or repair parts and then perform finishing actions such as painting the repaired portions of the vehicle, the prices for the parts listed, the prices for any materials listed and the prices for any labour listed).

The automated assessment platform 1310 of the example embodiment has three layers: an ingestion layer 1345; an assessment engine layer 1365; and an application layer 1380. In example embodiments, the automated assessment platform 1310 can assess the claim input through these layers 1345, 1365, 1380 in sequence.

In the example embodiment, the ingestion layer 1345 handles the input of data into the automated assessment platform 1320. The ingestion layer 1345 includes a standard application programming interface (API) 1335 and a secure file transfer protocol upload module 1340. The standard API 1335 provides a mechanism to input the claim input data 1305 into the platform 1310, in this embodiment by being connected to an in communication with the insurer systems into which the data has been input by a combination of the insured party, the garage and the insurer personnel. In the example embodiment, the insured party is

provided with a web application on their mobile phone to take photographs of the damage to their vehicle by the insurer and the image data captured using this web application is stored on the insurer computer systems and provided by the standard API 1335 to the assessment platform 1310. Further, in the example embodiment, the body shop estimate 1330 are
5 provided to the insurer by the garage using standardised third party software that captures the inputs from the garage personnel and prepares the body shop estimate 1330 in a standardised format for the insurer and uploads this directly into the insurer computer systems using the insurer's computer system API, which in turn provides this data to the assessment platform 1310 using the standard API 1335.

10 In the example embodiment, the assessment engine layer 1365 comprises a scope triage module 1350, a core assessment module 1355 and a workflow triage module 1360. More detail on the assessment engine later 1365 will be provided below with reference to Figure 14.

15 In the example embodiment, the application layer 1380 comprises a loss adjuster application 1370 and a results API 1375. The loss adjuster application 1370 allows the insurer personnel to review the detailed independent repair estimate that has been generated by the assessment engine layer 1365 against the supplied garage repair estimate 1120, 1330 to enable the insurer personnel to compare the details of each estimate, and to assist with efficient comparison the loss adjusted application 1370 highlights any differences between the
20 two estimates and provides an indication of the potential problems identified in the garage provided body shop estimate 1330. For example, the body shop estimate 1330 might list an additional labour operation that has not been determined to be required by the independent repair estimate generated by the assessment engine layer 1365, and this would be highlighted via the loss adjuster application 1370. The Results API 1375 allows for the output of the
25 assessment engine layer 1365 to be sent to the relevant insurer computer systems and personnel, for example to enable approval of a repair estimate 1330 in the case where the independent repair estimate generated by the assessment engine layer 1365 concurs or matches within a predetermined approval threshold set by the insurer personnel.

30 In an alternative embodiment, the supplied garage repair estimate 1120, 1330 is used as an additional input by the assessment engine layer 1365 which then outputs one or more determinations of incorrect, potentially fraudulent or inconsistent portions of the supplied garage repair estimate 1120, 1330.

35 In the example embodiment, in order to output an updated claim 1315 that can be or is approved by the insurer, the independently generated repair estimate from the assessment engine layer is output along with the claim input data 1305 and any result output as well as any adjustments made via the loss adjuster application 1330.

In some embodiments, the automated review (or assessment platform) 1310 can include, however is not limited to, one or more layers, modules, platforms, interfaces and/or tools. For example, the automated review may incorporate a platform 1310 to perform comprehensive automated assessments of claims using models, standard APIs to integrate with claims systems in the claim lifecycle and tools to help stakeholders use these assessments in their workflows. In some embodiments, the assessment platform 1310 can have two layers: an ingestion layer 1345 and an assessment engine layer 1365. In some embodiments, the ingestion layer 1345 can comprise a standard Application Programming Interface (API) 1335 and an SSH File Transfer Protocol (SFTP) upload 1340 or any other network protocol that may be used for secure file transfers. In some embodiments, the assessment engine layer 1365 can comprise a scope triage module 1350, a core assessment module 1355 and a workflow triage module 1360.

Figure 14 shows a more detailed diagram 1400 of the assessment engine layer 1365 shown in Figure 13 of the example embodiment which will now be described in more detail.

The claim data that is provided is assessed automatically using the automated assessment platform 1310. As mentioned above, in the example embodiment, the assessment engine layer 1365 of the automated assessment platform 1310 comprises three major modules: the scope triage module 1350, the core assessment module 1355 and the workflow triage module 1360.

In the example embodiment, the scope triage module 1350 performs data integrity checks 1405 on the data received from the ingestion layer 1345 by assessing claim photo coverage 1410, damaged panels photo coverage 1415, and vehicle identification number (VIN) availability and validity 1420.

The assessment of claim photo coverage 1410 determines, for the photographs 1320 that have been provided to the insurer, what part/region of the vehicle is shown in each photograph. This allows an assessment of whether sufficient photos have been taken of the vehicle.

The assessment of damaged panels photo coverage 1415 then performs verification against the order sheet 1325 and the body shop estimate 1330 that the photographs 1320 show all of the visible damage for which repairs, or replacements, are necessary in order for the automated assessment platform 1310 to make an independent assessment of the damage to the vehicle and the likely repair work required using the photographs 1320 in order to input these into estimatics platforms to generate an estimate. In some embodiments, if a predetermined threshold of coverage of the vehicle with the photographs 1320 provided is not met, then the process can terminate and a request be provided to provide additional images of the damaged vehicle, optionally indicating the parts/regions for which images are required.

In other embodiments, the process continues and an assessment is made by the system 1400 for the coverage possible with the images 1320 provided.

The determination of VIN (vehicle identification number) availability and validity 1420 checks that data is available for lookup for the supplied VIN and that the VIN matches either the detected or provided features of the vehicle.

In the example embodiment, the core assessment module 1355 is used to create an independent estimate of the repair work needed for the damaged vehicle based on the provided photographs 1320 that can then be used to verify the provided order sheet 1325 and body shop estimate 1330. The independent estimate of the repair work needed is generated by the core assessment module 1355 and checked using a series of checking modules 1430, 1435, 1440, 1445, 1450, 1455, 1460, 1465, 1470.

In the example embodiment, a leakage check module 1430 performs a check to identify any aspects of the body shop estimate 1330 that seems either unnecessary or superfluous; or to identify any decisions of the body shop estimate 1330 that seem to be incorrect (e.g. an incorrect repair versus replace decision) using data from the insurer's records of previous approved and rejected claims. The leakage check module 1430 in the present embodiment uses a repair replace check module 1435, a labour hours check module 1440, a paint check module 1445 and a parts check module 1450 to check respective aspects of the order sheet 1325 and the body shop estimate 1330.

The repair replace module 1435 checks that the repair and replace (i.e. that parts needing replacement and the parts needing repair rather than replacement) aspects of the order sheet 1325 and the body shop estimate 1330 match the equivalent aspects of the generated independent estimate of the repair work needed for the damaged vehicle.

The labour hours module 1440 checks that the labour (i.e. the amount of labour time and the labour operations/tasks to be carried out) aspects of the order sheet 1325 and the body shop estimate 1330 match the equivalent aspects of the generated independent estimate of the repair work needed for the damaged vehicle.

The paint module 1445 checks that the paint (i.e. the paint operations and materials cost) aspects of the order sheet 1325 and the body shop estimate 1330 match the equivalent aspects of the generated independent estimate of the repair work needed for the damaged vehicle.

The parts module 1450 checks that the parts (i.e. the list of parts needed) aspects of the order sheet 1325 and the body shop estimate 1330 match the equivalent aspects of the generated independent estimate of the repair work needed for the damaged vehicle and that these aspects are priced in line with data obtained from live data sources that can be queried for part prices.

In the example embodiment, a fraud check module 1455 performs a check to identify any aspects of the order sheet 1325 or the body shop estimate 1330 that seem fraudulent. The fraud check module 1455 in the example embodiment uses an inconsistent claims check module 1460, an inconsistent damage check module 1465 and a vehicle ID check module 1470.

The inconsistent claims check module 1460 checks for whether the order sheet 1325 or the body shop estimate 1330 list repair work that has not been predicted from the input data in the generated independent estimate of the repair work needed for the damaged vehicle. For example, a claim might be inconsistent if the reported point of impact is at the front of a vehicle then small paint scratches to the rear of the vehicle are unlikely to have been caused by the impact and are more likely to have already been present on the vehicle as damage prior to the impact.

The inconsistent damage check module 465 checks for whether the order sheet 325 or the body shop estimate 330 list repair work that is excessive compared to the repair work predicted to be necessary from the input data in the generated independent estimate of the repair work needed for the damaged vehicle. For example, damage might be inconsistent if a claim is made that a car had accidental damage occur overnight when it was parked but the assessed or reported damage has most or all of the characteristics of a vehicle that has been crashed by its driver.

The vehicle ID check module 1470 checks whether the VIN indicates the vehicle is the insured vehicle, that the vehicle is the correct make and model, and that the parts proposed to be used in the body shop estimate 1330 are consistent with the make and model of vehicle retrieved from looking up the VIN in a third party database.

In some embodiments, for example, the checks may be implemented to check against an independently generated repair estimate within higher and lower estimate thresholds for each of the modules 1430, 1435, 1440, 1445, 1450, 1455, 1460, 1465, 1470. In other embodiments, only some of the check modules 1430, 1435, 1440, 1445, 1450, 1455, 1460, 1465, 1470 are used to perform checks on the order sheet 1325 and body shop estimate 1330 against the independently generated repair estimate.

In the example embodiment, following the core assessment module 1355, the workflow triage module 1360 is then used to determine the output of the assessment engine layer 1365 or to determine what workflow is triggered as a result of the output of the assessment engine layer 1365.

The workflow triage module (or phase) 1460 in the example embodiment include assessment of usability of the output 1480, assessment of the value analysed 1485, determining a claim priority score 1490, and the application of any custom workflow rules 1495.

According to the example embodiment, the assessment of usability of the output 1480 is performed to assess whether the extent to which the output of the core assessment module (or phase) 1355 is of use. This assessment is performed using two modules: the value analysed assessment module 1485 and the claim priority score module 1490.

5 According to the example embodiment, the value analysed assessment module 1485 assesses what proportion of the body shop estimate 1330 have been checked against the independently generated repair estimate, as due to the use of photos to generate the independent repair estimate it is not possible to check all aspects of the repairs proposed to, for example, internal structural damage not captured or detectable from photos of the
10 damaged vehicle.

According to the example embodiment, the claim priority score module 1490 assigns a priority to the claim based on the level of leakage determined 1430 and the level of fraud determined 1455. This can be determined using a predetermined scoring system or a relative scoring system versus other claims being processed substantially concurrently.

15 Further, according to the example embodiment, there is application of any custom workflow rules 1495 that have been preconfigured by the insurer in order for the system to work with their internal policies, procedures and systems.

In other embodiments, the output of the core assessment module 1355 is used to trigger certain workflows or actions. For example, should there be no determined incorrect, potentially fraudulent, or inconsistent portions of the input claim then the input claim is automatically approved. In other embodiments automatic approval is made if only a portion of the input claim under a predetermined threshold amount (for example 20%, or 10% or 1% depending on insurer preferences) that is determined to be incorrect, potentially fraudulent, or inconsistent. In some embodiments, if an incorrect VIN is detected in an input claim then a
20 workflow can be triggered to avoid manual review of the claim and return the claim to enable the repair shop with a prompt to resubmit the claim with the correct VIN data. In other embodiments, where insufficient data is determined to have been supplied by the repair shop (e.g. no VIN, a lack of photos of the vehicle or of one or more portions of the vehicle) then a workflow can be triggered to again avoid manual (human) review and require the repair shop
25 to submit the claim again with the determined missing data. These automatically triggered workflows or actions can improve efficiency in processing claims by avoiding the need for manual (human) review.

In some embodiments, incorporating the described automated review process can improve process of assessing proposed repair work in a variety of ways including improved
35 damage assessment, the automation of damage assessment, the more efficient and/or prompt authorisation of repair work, and faster and more efficient settlement assessment. Damage can be assessed more efficiently and accurately by triaging claims that need to be processed

in different ways, based on for example the value of the claim that can be analysed by the system automatically and/or the priority of the claim. The use of automated assessment can accelerate finding leakage in claims while allowing the quick approve claims with no anomalies automatically. The automated review process of other embodiments can be used to assess and flag post-repair leakage by reviewing approved invoices and performing post repair audits. In some embodiments, the various assessments can be used for triaging claims to determine whether the best way to process a claim is to send an expert on-site if structural damage has been identified, partial processing to quickly identify leakage or fraud on a claim for insurer personnel to review and challenge, and automation to automate approval of low damage claims or high certainty claims based on insurer predetermined statistical or absolute thresholds set for damage or leakage for example.

Figure 15 shows a comparison analysis process 1500 as implemented in one example embodiment where the process is used to perform the checking of the images of the damage to the vehicle and the repair estimate submitted by the repair shop using a model trained on historical data.

Specifically, Figure 15 shows a flowchart depicting a method 1500 for assessing proposed vehicle repair estimate data against historical data and predetermined approval criteria by processing image input 1505 of the damage to the vehicle to be repaired using an image classifier 1510. In this embodiment, the vehicle repair estimate data is analysed using a (decision) classifier 1515, which has been trained using historical vehicle repair data and works in conjunction with an estimatics platform 1535 by determining the damage to the vehicle and uses the estimatics platform to determine repair data for the determined damage in order to verify the input vehicle estimate data 1520, to output any identified errors or anomalies 1525 in the input vehicle estimate data 1520. In some embodiments, the input vehicle estimate data 1520 along with any errors or anomalies 1525 can be displayed to a user in a user interface on a computer for review, for example in a loss adjuster application.

Figure 16 shows a comparison analysis process 1600 as implemented in another example embodiment where the process is used to perform the checking of the images of the damage of the vehicle against historical data by generating an independent repair estimate 1615 and comparing this to the submitted repair estimate 1620 from the repair shop.

Specifically, Figure 16 shows a flowchart depicting a method 1600 for assessing proposed vehicle repair estimate data 1620 against historical data and predetermined approval criteria by processing images of the damage to the vehicle 1605 using an image classifier 1510. In this example embodiment, an independent repair estimate is generated by the (estimate) classifier 1615 in conjunction with lookups made to an estimatics platform 1635, where the image and classifiers 1610, 1615 have been trained using historical vehicle repair data and work in conjunction with an estimatics platform 1635, to output an independent repair

estimate. The independent repair estimate and the repair estimate 1620 provided by the repair shop are compared 1625 to identify errors or anomalies 1630 in the input vehicle estimate data 1620. In some embodiments, the independent repair estimate and input vehicle estimate data 1620 along with any errors or anomalies 1630 can be displayed side-by-side to a user in a user interface on a computer for review, for example in a loss adjuster application.

Example embodiments, such as the example embodiment process 1700 shown in Figure 17, can be used by body shops to enable them to check their proposed repair estimate and, if any irregular operations are identified, change their proposed repair estimate to remove any irregular operations or improve the supporting evidence for necessary operations flagged as irregular before the repair estimate is sent to the insurer, in order to comply with the requirements of automated assessment of their repair estimate by the insurer using the automated assessment platform.

By previewing the claim estimate, and the associated automated assessment of the claim estimate, the body shop can be made aware of potentially incorrect decisions used to generate their repair estimate or whether they will be required to submit more information (such as photographs to support their proposed repairs, for example). In some embodiments, the system/platform can be used to produce an estimate automatically for the body shop by analysing photographic images of the damaged vehicle obtained by the body shop and obtaining repair data using estimatics providers.

In the embodiment shown in Figure 17, a process for pre-checking a repair estimate 1700 is shown and will now be described. When an accident occurs 1705, the client reports the accident and details of the accident and the damage to the insurer 1710. The client takes their damaged vehicle to a garage for repair 1715. The garage surveys the damage and takes photographs of the damaged vehicle, and inputs these photographs and some details of the vehicle and damage/accident into the estimate platform 1720. The estimate platform used a model trained on car damage data to prepare a determination of the damage to the vehicle and the severity of this damage, including for example whether to replace or repair parts that have been damaged. This determination is submitted to an estimatics platform to determine for example the parts list required and the labour hours and costs in order to prepare the repair estimate. The repair estimate is then provided to the insurer 1725. Since the platform has already prepared the repair estimate using the verified process for assessing the damage to the vehicle, the repair estimate can either be pre-approved or can be verified by the insurer software in the way used for manually prepared repair estimates. Once approved, the garage can begin work.

Optionally, should the estimate have been prepared incorrectly, or a warning message is overridden or ignored by the user and so a claim is submitted that isn't pre-approved, it is

possible for the insurer to process the estimate 1730 and reject it, returning the process to the step of the garage preparing an estimate 1720 for resubmission 1725.

Referring now to Figures 18 to 20, several embodiments illustrating one or more further aspects will now be described.

5 Referring first to Figure 18, a generalised embodiment 1800 will now be described in more detail to provide an overview of the approach taken by the methods of various aspect/embodiments.

10 Images 1805 are captured of a damaged vehicle, typically these are a plurality of images of the vehicle from a variety of viewpoints of the vehicle and typically these show the damage to the vehicle both close-up and also in context with undamaged portions of the vehicle. These images are typically captured as part of the process of making an insurance claim, but may also be captured by an occupant or owner of the vehicle shortly after the damage occurs or by a vehicle repair business as part of documenting their damage assessment.

15 The images 1805 are provided to a plurality of image classifiers 1815, each classifier 1815 outputting a decision on whether each of the images 1805 is classified as containing any damage to a specific normalised portion of the vehicle (e.g. the front wing; the front left door; the rear bumper; etc). In other embodiments, a single classifier 1815 may be used to determine damage classifications to more than one portion of the vehicle or the entire vehicle.

20 The output of the image classifiers 1815 can be used a set of decisions (or in some embodiments classifications, optionally with confidence values), each being a decision per portion of the vehicle, as to whether there is damage to that portion of the vehicle. In some embodiments, a level of damage is also indicated in the decision that is output from the image classifiers 1815 (e.g. minor or major; or damage categories such as retouch, repair or replace).
25 In some embodiments, one or more locations of damage are indicated in the decision that is output from the image classifiers 1815. In some embodiments, other features are also output from the image classifiers 1815.

In addition to the images 1805, a claim input 1810 is provided in the form of an electronically generated detailed list of parts and labour operations to perform a repair to the
30 damaged vehicle. The claim input 1810 is previously generated using a third party database, provided with inputs by vehicle repair personnel based on their assessment of the damage to the vehicle, which generates an output list of parts for the specific model of vehicle and repair operations required to repair or replace each of the damaged components of the vehicle for the vehicle repair personnel.

35 The claim input 1810 is provided to a NLP (natural language processing) model 1825 which reads the human readable text in the claim model 1810 and converts the claim input 1810 into normalised data for each portion of the vehicle (for example, converting a portion of

the claim input comprising a detailed list of ten parts and corresponding list of manual operations and associated labour times into normalised data such as “replace front left wing, three hours total labour time, requires X parts” where X could be a financial value, a list of specific parts or a number of parts). In some embodiments, any associated photos from the claim input 1810 are included in the normalised claim input generated by the NLP model 1825. In embodiments, the NLP model 1825 comprises a plurality of NLP models.

Natural Language Processing (NLP) is a field of machine learning where a computer is able to understand, analyse, manipulate and generate human language. The NLP model 1825 uses techniques from this field to understand the human readable text in the claim model 1810 and convert this into a standardised set of information that is used to train and at runtime with the machine learned models of the embodiment. Building a mapping and parsing system using NLP can allow the method of this and other embodiments to correctly ingest estimates from third party estimatics systems, being known or new estimatics systems, with zero training data, even in new markets/jurisdictions. In some embodiments, a character-level language model is used to predict based on a part description which can then be mapped to a normalised portion of the vehicle (or normalised part), for example a front bumper. In embodiments, the language model is trained on historical unlabelled estimate data and/or across different languages – this can make the learning process more efficient and/or increase performance at inference and/or allow the model(s) to work better when applied to new geographies/markets/jurisdictions.

The decision(s)/classification(s) from the image classifiers 1815 and the normalised claim input from the NLP model(s) 1825 are provided to a classifier 1820. The classifier is a trained model or models that assesses the output from the NLP model(s) 1825, which represents the repair estimate generated by the vehicle repair personnel, to the output from the image classifier(s) 1815, which represents an independently generated damage estimate (and in some embodiments the method can use one or more third party databases to generate a detailed list of parts and labour operations based on the independent generated damage estimate output from the image classifier(s) 1815), in order to classify whether the claim input 1810 (in whole or in part) is correct.

The output of the classifier(s) 1820 can be used to verify a repair estimate generated by the vehicle repair personnel, for example by the vehicle repair personnel before it is submitted for approval by a customer or an insurer (to ensure it is correct) or by an insurer (to ensure it is correct to approve or reject a claim for the correct repair work from an insured customer, for example as part of an insurance claim approval process).

To provide more detail on the concept of normalised car parts 1900, as used in above described example embodiment, the following detail is provided with reference to Figure 19.

In Figure 19, there is shown a generic model of an example vehicle 1910. The example vehicle 1910 shows a selection of normalised parts/regions 1900 that are common to most if not all vehicles, which common normalised parts 1900 can be used to describe in generic terms the typical parts found of most vehicles that might need to be considered by the method
5 for assessing or verifying estimates for repair work.

For example, the example vehicle 1910 has shown a front right wing 1920, a front left wing 1930, a left front door 1940 and a left back door 1950. Other possible normalised vehicle parts 1900 might include (non-exhaustively): bumpers, wings, wheel arches, roofs, windows, wheels, lights, grills, vents, radiators, mirrors, windscreens, and convertible roofs.

10 Using this normalised part schema allows the conversion of detailed repair estimates provided by vehicle repair personnel into normalised repair inputs, for example to convert a detailed part list and detailed labour times and operations into a normalised indication, for each normalised portion of the vehicle, of what damage is to the vehicle of interest and a summary of the parts and labour required to repair each normalised portion of the vehicle.

15 Similarity, using the normalised part schema, an analysis can be performed on the image data of the vehicle to assess which of the normalised parts of the vehicle exhibit damage in the photos of the vehicle and, in some embodiments, the extent (e.g. entire part/majority of part/minority of part) and/or severity (e.g. high/medium/low) and/or classification (e.g. replace/repair/paint) and/or location (e.g. position or positions on the vehicle) and/or other
20 features of the damage to each normalised portion.

With reference now to Figure 20, a more detailed example embodiment of the independent damage generation process 2000 will now be described more fully.

The input to the process are a plurality of images 2005, which comprises a set of images of the damaged vehicle from different viewpoints and including image data of the
25 vehicle as a whole and image data pertaining to the damage to the vehicle.

The images 2005 are provided to a Visual AI model 2015, in this example embodiment a machine learned model trained on historic appraiser decision data which includes image data of the damaged vehicles and the corresponding appraiser decisions for the damage on these damaged vehicles. Other machine learned and artificial intelligence approaches can be
30 employed in other embodiments to provide a Visual AI model 2015. The Visual AI model 2015 outputs a damage mask 2010 and a damage classification for each normalised portion of the damaged vehicle.

In some embodiments, using the normalised part schema/arrangement, and determining the extent, severity, location and type of damage can allow for the use of models
35 or look-up tables (e.g. using third party databases or by generating look-up tables) to determine a more detailed level of granularity in output determinations (e.g. specifying the exact part number for a damage bumper for the make, model and year of vehicle being

considered and how long it will take to replace the part). Thus the benefit of a normalised part schema can be that it allows a generalised visual assessment across many types of vehicles and/or geographies (allowing the visual modelling to be performed more easily and/or robustly) without losing the ability to determine a detailed output (e.g. specific parts and/or operations and/or labour time requirement).

The damage mask 2010 comprises a damage classification for all normalised parts of the damaged vehicle.

The damage classification for each normalised portion of the damaged vehicle is output from the Visual AI model to the repair model 2020.

The repair model 2020 comprises a set of models including a set of repair/replace models and a repair labour hours model, all of which are machine learned models trained on historic appraiser decision data which includes image data of the damaged vehicles and the corresponding appraiser decisions for the damage to these damaged vehicles. Other machine learned and artificial intelligence approaches can be employed in other embodiments to provide a repair model 2020 and/or repair/replace models and/or one or more repair labour hours models. The output of the repair model 2020 is a classification, for each of the normalised parts of the vehicle, of whether the normalised portion of the vehicle requires repairing or replacing and the labour operation category (e.g. minimal/straightforward/complex or a rough estimate of time required).

A blend engine 2025 receives both the damage mask 2010 and the output of the repair model 2020. The blend engine 2025 then assesses whether any paint blending will be required for each of the normalised parts of the vehicle. Paint blending is the process of ensuring a smooth transition in paint colour between neighbouring parts of a vehicle, thus may be required (for either a repaired/replaced part and/or at one or more neighbouring undamaged parts) to ensure a contiguous finish to the paint colour if a pre-painted replacement part is used to replace a damaged part or if a damaged part is painted following repair or replacement.

The output of the blend engine 2025 is provided to the paint rules engine 2030 and the remove and install (R&I) rules engine 2035.

The paint rules engine 2035 comprises a set of rules used to assess the output of the blend engine and the damage mask 2010 to estimate the likely materials and labour required to paint the repaired vehicle. The set of rules in the example embodiment are hand crafted but in other embodiments a machine learned approach can be used.

The R&I rules engine 2035 determines which parts need to be removed from the vehicle so that the physical repair operations can be carried out and then re-fitted. In some embodiments, the output from the R&I rules engine 2035 is an aggregate amount of time that needs to be spent on removing these parts from the vehicle and re-fitting these parts once repairs have been made, per normalised part of the vehicle.

Based on the outputs of the Visual AI model 2015, the repair model 2020, the blend engine 2025, the paint rules engine 2030 and the R&I rules engine a database lookup 2040 is performed at a 3rd party databased 2050 to generate an electronically generated detailed list of parts and labour operations to perform a repair to the damaged vehicle, using the assessment generated of the damage to the vehicle and details about the vehicle (e.g. make, model, year, etc if any or all of this data is available).

The output of the damage assessment process 2000, once the detailed list has been prepared using the database lookup 2040, is reviewed by a set of overlap rules at an overlap rules engine 2060 to determine if any duplicate parts or labour operations are listed in the detailed list, which can then be removed from the detailed list by the overlap rules engine 360 and an output produced 2070.

The output 2070 can then be compared to an electronically generated detailed list of parts and labour operations to perform a repair to the damaged vehicle prepared by vehicle repair personnel.

By breaking up the process of generating a repair estimate into human decisions and software-assisted decisions, and building machine learned models to predict the human decisions, a completely automated approach to generating a repair estimate can be adopted.

In some embodiments, some of the models used can have a limited capacity to prevent overfitting to the training data, for example by applying a rule-based system or using a shallow, capacity constrained machine learning model/approach.

In some embodiments, one or more of the models are trained in a manner agnostic to the make, model and year of the vehicles being considered. This can allow the training of a universal set of one or more models, or even a universal approach, to generating classifications of damage/damage estimates for vehicles or a generalisable damage representation (from, for example, just image data of the damaged vehicle). In embodiments, some or all of the models comprise a final shallow layer which takes the generalisable damage representation and adds the make, model and year information for the vehicle being considered in order to make the final prediction/classification. In these embodiments using the final shallow layer to add the make, model and year information, the approach can require much less training data to train the model(s) than a full convolutional neural network.

In some embodiments, the model(s) are trained using data from multiple geographies, such that the first layers of the neural networks in the model(s) processes the input data in a universal way across all geographies in order to extract generalisable features for vehicles, then each geography is refined using only training data from that geography in order to learn the specific repair methodologies for each geography. In embodiments, the training dataset can be pseudo-labelled with geography specific labels so that the geography specific branches of the model can be treated as branches in a multi-task learning technique. In other

embodiments, one or more models can be trained to re-weigh the loss function to account for the proportion of repair/replace decisions in each geography. In other embodiments, one or more models can additionally have a confusion loss layer present in substantially the middle of the network, which can be used to enforce that at that point in the network it doesn't
5 differentiate between images coming from different geographies.

In some embodiments, the method normalises the point or impact (e.g. the front left of the vehicle) and/or the object involved in the collision (e.g. a tree, or another car) and includes a further step of checking that the damage prediction/classification is consistent with either or both of the point of impact and the object involved in the collision.

10 An alternative embodiment will now be described in more detail where no previous data has been acquired in respect of a new situation, for example where a new make, model and year of a vehicle is encountered.

In this embodiment, upon encountering a new make, model and year of a vehicle, the system reverts to using a "default" vehicle type (or, in other embodiments, the system may try
15 to determine a generic car type such as a "sedan", "pick-up" or "van" and use pre-determined or pre-learned data for that vehicle type). The system uses either the "default" or the selected "generic" vehicle type at every subsequent step where otherwise a specific make, model and year is used (for example as input to a final layer of the visual models, or as an input to the blend engine and/or paint rules and/or R&I rules and as input when querying the third party
20 database).

By training a generic model that uses normalised parts per vehicle, new models of vehicle can be analysed immediately (using the new make/model as an input) and the machine learned models can learn correlations with existing learned parameters for other make/models. In this embodiment, a default setting for the make and model of a vehicle is
25 used, having been trained on training data lacking make, model or year information, allowing "zero-shot" prediction for the damage representation for a vehicle of unknown make, model and year (but with an expected lower accuracy than if at least some training had been performed specifically for the make, model and year).

Another embodiment will now be described in more detail where a trained approach in
30 one jurisdiction is adapted to a different unseen jurisdiction.

In such an embodiment, the system is configured as in the previous embodiment described in relation to Figure 20, but the design decisions made allow the method to be adapted to a new geography with little data. A relatively small amount of data is collected in respect of the different unseen jurisdiction (i.e. new data for the new jurisdiction). In particular,
35 by finetuning just the last layers of the visual models (which requires a lot less data than training a new visual model, hence the small amount of data collected for the new jurisdiction), using this new data and by adjusting the blend, paint and R&I engines based on the new data

(which also requires much less data than training machine learning models from scratch to perform these operations), the previously-described approach(es) can be adapted to a new domain (i.e. to the different unseen jurisdiction).

By using domain adaption, a model can be trained to work universally but then adapted to be used with different markets/jurisdictions/models. In some embodiments, a model is generated for one market (for example, the insurance repairs market in France) and then, using domain adaption, fine-tuned for use in another market (for example, the insurance repairs market in Japan). For example, the initial dataset used to train the models may have a very large amount of data compared to the following dataset (using the previous example, there might be one hundred times the data that can be used to train the model for France than for Japan). The differences between markets/jurisdictions/models may be any number of factors, for example in some jurisdictions it is typical to replace parts rather than repair them. This approach can also allow a high level of performance when adapting to local changes (e.g. a hailstorm or other natural disaster has caused widespread damage) by adjusting the prior distribution of labels. In some embodiments of this approach, live key metrics are tracked, including for example metrics such as the proportion of normalised parts that are repaired vs those which are replaced. When a drastic shift in distribution of tracked metrics occurs, it would indicate something significant has changed in how vehicles are getting damaged and/or are being repaired or proposed to be repaired (for example, a hail storm or natural disaster). Using thresholding, the approach embodied in the system can be adjusted accordingly (i.e. when the distribution has changed by over a predetermined amount) by tuning any one or more of the blend/paint/R&I engines or by fine-tuning the top layers of the visual classifiers in order to allow adaption to changes in repair methodologies quickly.

In some embodiments, an efficient categorical feature encoding is used for the make and model in order to further reduce the amount of training data needed.

Referring now to Figure 21, an alternative embodiment implementing a more detailed set of normalised vehicle parts 2100 (or normalised/standardised vehicle parts or regions) will now be described in more detail. This more detailed set of normalised vehicle parts 2100 can be used with any of the other described aspects and/or embodiments herein.

Shown in Figure 21 is a generic vehicle diagram according to an embodiment, where the parts/regions of the vehicle are broken up into a plurality of normalised/standardised parts, zones or regions 2100. In Figure 21, there are shown the following normalised/standardised parts: a front left headlight 2102; a front radiator grille 2104; a front right headlight 2106; a front bonnet 2108; a front right fender 2110; a front right wheel 2112; a right running board 2114; a front right door 2116; a front right door window 2118; a back right door 2120, a back right door window 2122; a trunk 2124; a rear right quarter panel 2126; a rear right light 2128; a rear bumper 2130; a rear left light 2132; a rear left quarter panel 2134; a rear window 2136;

a back right door window 2138; a back right door 2140; a left running board 2142; a front left door 2144; a front left door window 2146; a front window 2148; a front left wheel 2150; a front left fender 2152; and a front bumper 2154. In other embodiments only some of these normalised parts may be used and/or some other normalised parts may be exchanged for those used in the example embodiment.

Referring now to Figure 22, a paint check process 2200 according to an example embodiment is shown and will now be described in more detail.

In this embodiment, the input data is a plurality of photos 2202 of the damaged vehicle and a repair shop estimate 2204 which includes details of the proposed repair operations and/or materials required to repair the damage to the damaged vehicle (typically along with the proposed cost for the proposed repair operations and/or materials).

Initially, some pre-check checks 2206 are performed. The pre-check checks 2206 include (a) determining whether the claim is eligible, i.e. determining whether the proposed repair operations and/or materials are covered by an insurance policy; (b) determining whether the proposed repair operations including proposed painting of an in-scope and/or out-of-scope panels of the damaged vehicle; and (c) to determine whether there are look-up tables for the specific make, model and year of the damaged vehicle in order for the system to be able to retrieve the relevant data required to determine the paint requirements for the damaged vehicle. In other embodiments, only some of these checks are performed and/or alternative or modified checks may be performed.

Following the pre-check checks 2206, the paint area is determined 2208. Specifically, the areas of the damaged vehicle that require painting are determined. To do this, per-part models are used to determine whether each of the parts of the vehicle require painting and to what extent – i.e. does only a spot on a panel need to be painted, or does a portion of the panel need to be painted, or does most/all of the panel need to be painted. This is described in more detail below in relation to Figure 23 and the underlying models used in the process 2200 can be adapted to different jurisdictions by for example applying different training weights. Adapting models to new jurisdictions can be done in various ways. In one embodiment, it can be done by using domain adaption methods with data from the new jurisdiction. In another embodiment, it can be done by adjusting the weight of each class (e.g. spot vs minor vs major, as described elsewhere) based on the prevalence in that jurisdiction. In a further embodiment, it can be done by setting different thresholds on the output of the classifier to determine what level of damage corresponds to each class (e.g. spot vs minor vs major, as described elsewhere).

Once the paint area has been determined, the paint time is determined 2210 – specifically an estimate for the time it will take to paint the determined paint area is generated. To do this, in this embodiment, a look up is performed (i.e. a query sent) to a third party

database is performed for all of the parts/panels determined to require painting, including details of the extent to which each of the panels require painting, and the third party database returns a time value based on the query, typically the time value based on the specific make, model and year of vehicle.

5 In parallel with the paint area determination process 2208, a pre-painted panel check 2212 is performed. The pre-painted panel check 2212 determines whether the damage to the damaged vehicle will require any pre-painted panels to be used to repair the damaged vehicle, thus removing the need to paint one or more parts of the vehicle due to the parts used to repair the vehicle being pre-painted.

10 For the pre-painted panels determined to be required by the pre-painted panel check 2212, two further processes are performed: a determination of the replace paint time 2214; and a determination of the addition cost 2216.

The determination of the replace paint time process 2214 comprises determining the time required to perform the replacement of the damaged panel(s) with the pre-painted panel(s). To do this, in this embodiment, a look up is performed (i.e. one or more queries are sent) to a third party database for all of the parts/panels determined to require replacement with a pre-painted panel, and the third party database returns a time value based on the query, typically the time value based on the specific make, model and year of vehicle.

15 The determination of the addition cost process 2216 comprises determining the requires auxiliary task that need to be performed in order to replace any parts with pre-painted parts or prepare the damaged vehicle for painting the areas determined to require painting, including: determining the time required to apply masking tape to prevent paint coating areas of the car not to be painted; and to prepare the paint to be applied (e.g. to mix the paint).

A total paint time is then determined by combining the determined repair paint time from the paint time determination process 2210, the addition cost process 2216 and the determine replace paint time process 2214. In some embodiments, overlap between the times can be subtracted from the total time determined by the process 2218.

20 Following this, a paint labour cost is determined 2220, using the total time determined 2218 and a lookup performed using a third-party database to obtain appropriate labour rates for the time required to perform the painting operations.

25 Following this, a material cost is determined 2222. Specifically, the cost of the paint required for the paint areas 2208 is determined from by adding together the materials cost of the paint and the pre-painted panels required. The cost of the materials can be determined for the make, model and year of the damaged vehicle by performing a third party database lookup to retrieve the values required.

35 A total paint cost is then determined 2224 by combining the labour cost 2220 and the material cost 2224.

Then some post-check checks 2226 are performed, including in this embodiment: checking that unnecessary parts of the vehicle aren't being proposed to be painted (e.g. the door frame or the inside of a panel), and/or including items inconsistently (e.g. as materials but not labour proposed), in the proposed repair operations and/or materials.

5 Finally, one or more decisions as to whether the proposed repair shop estimate 2204 contains any unnecessary or anomalous proposed repair operations and/or materials is generated based on the pre-check checks 2206, the post-check checks 2226 and a comparison of the determined total paint cost 2224, the determined material cost 2222, and the paint labour cost 2220 to the proposed repair operations and/or materials in the repair
10 shop estimate 2204.

As shown in Figure 23, it should be noted that different jurisdictions/geographies have different thresholds, preferences, schema and rules that impact what is acceptable in proposed repair operations and/or materials used to repair damaged vehicle – an example of potential differences between jurisdictions 2320 is shown in table 2300 and will now be
15 described in more detail below.

In table 2300, there are shown five generic example jurisdictions A, B, C, D and E. Each of the jurisdictions have different rules 2322, 2324, 2326, 2328, 2330 for what to do per assessment of paint requirements 2330. The "AI decision", i.e. the decision/classification of the per-part model assessing whether a part of a damaged vehicle requires repainting in this
20 example embodiment can output four classifications: none; spot; minor and major. However, to adapt these classifications to the relevant jurisdictions, the jurisdiction specific rules 2300 need to be selected as appropriate. So for example, in jurisdiction A the rules 2322 specify that a spot paint job requires a fifth of the panel to be painted while a minor paint job requires half of the panel to be painted and a major paint job requires the whole panel to be painted.
25 In contrast, in jurisdiction B the rules 2324 specify that a spot paint job requires just the spot of the panel to be painted while a minor paint job requires less than fifty percent of the panel to be painted and a major paint job requires more than fifty percent of panel to be painted. Further, in jurisdiction C the rules 2326 specify that a spot paint job requires just the spot of the panel to be painted while a minor paint job requires twenty percent or less of the panel to
30 be painted and a major paint job requires more than twenty percent of the panel to be painted. In further contrast, in jurisdiction D the rules 2328 specify that a spot paint job requires just the spot of the panel to be painted while a minor paint job requires the panel to be painted at a cost fifty percent below fully repainting the panel and/or replacing the panel with a pre-painted panel and a major paint job requires the panel to be painted at a cost twenty percent below
35 fully repainting the panel and/or replacing the panel with a pre-painted panel. Finally, in jurisdiction E the rules 2330 specify that a spot paint job requires thirty percent of the panel to

be painted while a minor paint job requires sixty percent of the panel to be painted and a major paint job requires ninety percent of the panel to be painted.

Referring now to Figure 24, to implement the models in the above embodiments, in an example embodiment, a model architecture 2400 is used including the use of a plurality of gradient booster trees 2440 which will now be described in more detail below.

Specifically, the model architecture 2400 used is a two-layer architecture using both visual models 2420 and gradient booster trees 2440.

The input to the visual models 2420 are the photos of the damaged vehicle 2410 and the repair shop estimate 2430. The visual models 2420 are trained on a per-part basis to determine, for each part of the damaged vehicle, whether and what degree of painting is required: none; spot; minor or major. These visual models are run on photos of all parts of the damaged vehicle (as shown in the photos/images). In addition, in this embodiment, visual models are also run to determine: (a) whether a replace or repair operation is required for each panel; (b) the amount of repair labour is required for each panel; (c) whether each part is damaged or undamaged. All of the scores from these visual models (painting; repair/replace; labour; damaged/undamaged) are input into the gradient booster trees 2440, across multiple parts, and the output is a prediction of a final score. The purpose of the gradient booster trees 2440 are to take information from the multiple visual models, and from more than one panel (e.g. just two up to all of the panels on the vehicle) and use this information to make a final paint determination in order to improve on just using a visual model 2420 to determine paint requirements per panel and combine the scores. The gradient booster tree 2440 comprises multiple decision trees having multiple nodes 2450a, b, c, d, e, f, g and multiple terminal nodes 2460a, b, c, d, e, f, g, h. Each individual decision tree can be quite a weak predictor by itself, but the boosting algorithm used can concentrate each one on a region of the prediction space such that, in aggregate, they form a strong predictor of the damage. In this embodiment, this is done by adding trees in an iterative fashion, using an algorithm such as AdaBoost, details of which are herein incorporated by reference.

Referring now to Figure 25, an equation for calculating paint cost 2500 according to an embodiment is presented and will now be described in more detail below.

The total paint cost 2510 is shown to comprise the sum of multiple functions including the labour rate 2520, the material cost rate 2530, the paint type and vehicle details 2540, the repair paint times 2560, the total repair paint time 2550, the total replace paint time 2570, the repair paint times 2580 and the addition paint time 2590.

Referring now to Figure 26, an overview of the damage determination process 2600 according to an embodiment will now be described in more detail below.

As an input to the damage determination process 2600, the photos of the damaged vehicle 2610 are provided. In this process, there is a need to determine if a normalised panel

is undamaged or if there is some damage present. The photos 2610 are provided to a visual damage classifier which provides classification of the damage to each part of the vehicle for each of the photos using classifiers trained per part that classify each of the photos 2610. Following this, a process of image segmentation is performed 2630 to identify segments of interest, or to segment damage within the images. Both the output of the visual damage classifier 2620 and the segmentation 2630, along with the photos 2610, are then used to output a final damage score 2650 using the damage classifier 2650. This can allow the final damage classifier 2640 to use the segmentation mask 2630 with the photos 2610 to better focus on and/or assess the segmented damage shown in the photos. In some embodiments, the segmentation 2630 outputs a segmentation mask indicating where it is most likely damage will be determined to have occurred from the photos 2610. In some embodiments, the visual damage classifier 2620 may output with very low certainty where damage has been determined in the photos 2610.

Referring now to Figure 27, an alternative embodiment of a damage determination process 2700 is shown and will now be described in more detail.

Specifically, this embodiment is concerned with damage such as cosmetic damage to a vehicle, which can be very hard to detect in images of the damaged vehicle, especially where for example the vehicle has a very light-coloured paint. Typically, if the images of the damaged vehicle are down-sampled during the inference process by the models, it is possible that such damage becomes difficult or impossible to determine. In the process 2700, the photos 2710 that are input are cropped 2715 in a random fashion to generate multiple cropped images for each image. In other embodiments, a part classifier can be used instead of performing random cropping in order to select particular portions of the vehicle that are likely to contain the part(s) of interest). The cropped images are then processed by the visual damage classifier 2720 to determine any damage on a per-part basis for each cropped image. The cropped images 2715 and the original images 2710 as well as the output of the visual damage classifier 2720 are provided to the image segmentation process 2730 as per the embodiment of Figure 26 and the damage classifier 2740 outputs a final damage score 2750.

Referring now to Figure 28, further details of the segmentation 2800 will now be shown and described in more detail.

The figure shows an example vehicle 2810 as might be shown in an image of a damaged vehicle, as well as a cropped portion of the image 2840 in which damage 2850 is more clearly visible and which can be segmented using a bounding box (or similar technique) 2860. Segmentation data can be collected, for example from human labellers creating training data, and a segmentation model trained on this data to perform segmentation of images (cropped or otherwise) to identify damage or regions of interest in the images.

In some situations, the images input will appear to show no damage in most of the images but only in the close-up images can the damage been seen. To output a substantially accurate score, the models need to be able to determine that the part is damaged in spite of the damage not being visible in most of the images of the part and in addition the model needs to be able to locate the damage visible in the close-up images to a location on the vehicle, which can sometimes be difficult as there is limited context to close-up images.

The segmentation performed on the image data can then be used by the damage classifiers to improve the ability to identify otherwise difficult to detect damage to vehicles, for example cosmetic damage that would otherwise only comprise of a very small portion of an uncropped image and which might only be visible in a close-up image. Thus, segmenting the area on the vehicle where the damage is located can assist the models to determine both that the damage is present and also where the damage is on the vehicle, as the damage might be identifiable in a close-up image of the vehicle but due to the close-up nature of the image the model may be unable to locate the damage on the vehicle.

Referring now to Figure 29, an overview of the damage determination process 2900 according to an embodiment will now be described in more detail below.

Multiple images 2910 are input into the process 2900, each of which is cropped into multiple crops of each image 2920', 2920'', 2920''', 2920ⁿ. Each of the multiple crops of each image 2920', 2920'', 2920''', 2920ⁿ are then processed by a multi-image model 2930. In this embodiment, the multi-image model 2930 comprises a visual damage classifier 2940, followed by a pooling operator 2950 then a final damage classifier 2960, which outputs a final damaged/undamaged score 2970 for each part.

By cropping the multiple images 2910, the image resolution of the images are preserved in the cropped images 2920', 2920'', 2920''', 2920ⁿ so that the full detail of the original resolution of the images can be processed by the multi-image model 2930 without needing to down-sample the images 2910. In alternative embodiments, instead of cropping the images 2910, the model 2930 can be trained to work at the original resolution of the input images without down-sampling.

By training and using a multi-image model 2930, which only outputs a prediction score of whether each part is damaged or undamaged 2970 after looking at all of the images 2910, the model 2930 can substantially overcome the limitations in the input data only showing damage clearly in some images of the vehicle and not in others. The multi-image model 2930 achieves this by the visual damage classifier 2940 producing a feature representation for each of the images 2910 and/or multiple cropped images 2920', 2920'', 2920''', 2920ⁿ and co-ordinate these representations using the pooling operator 2950 to concatenate the representations across all images and use a final damage classifier model 2960 to make final predictions. In alternative embodiments, the pooling operator 2950 can use a pooling

operation, e.g. a mean, max operation, with the representations produced by the visual damage classifier 2940.

In alternative embodiments, the cropping is optional and the multi-image model 2930 can be used with the original images 2910.

5 Referring now to Figure 30, which show an alternative embodiment 3000 to that shown in Figure 29, the alternative embodiment using a graph neural network, and this alternative embodiment will now be described in more detail.

Multiple images 3010 are input into the process 3000, each of which is cropped into multiple crops of each image 3020', 3020'', 3020''', 3020ⁿ. Each of the multiple crops of each
10 image 3020', 3020'', 3020''', 3020ⁿ are then processed by a multi-image graph model 3030. In this embodiment, the multi-image model 3030 comprises a visual damage classifier 3040, then a graph convolutional neural network 3060, which outputs a final damaged/undamaged score 3070 for each part.

As in the previous embodiment, in this embodiment by cropping the multiple images
15 3010, the image resolution of the images are preserved in the cropped images 3020', 3020'', 3020''', 3020ⁿ so that the full detail of the original resolution of the images can be processed by the multi-image model 3030 without needing to down-sample the images 3010. In alternative embodiments, instead of cropping the images 3010, the model 3030 can be trained to work at the original resolution of the input images without down-sampling.

20 By training and using a multi-image graph model 3030, all of the images 3010 can be considered when making a prediction 3070, as the graph neural network 3060 treats each feature representation as a node in a graph. The model 3030 can thus also substantially overcome the limitations in the input data only showing damage clearly in some images of the vehicle and not in others, as the model can have arbitrary connections between nodes in the
25 graph to represent this. The multi-image model 3030 achieves this by the visual damage classifier 3040 producing a feature representation for each of the images 3010 and/or multiple cropped images 3020', 3020'', 3020''', 3020ⁿ and then using the graph neural network 3060 to make final predictions.

Referring now to Figure 31, there is shown a process using a language model to extract
30 information relevant to visual damage 3100, according to an embodiment, which will now be described in more detail.

To determine whether damage is consistent with the reasons for the damage, one can compare the details of the accident/damage as reported (for example to an insurance company) and the determined damage to the vehicle.

35 To do this, in this embodiment, images of the damage to the vehicle 3110 are processed by a per-part visual model 3120 to generate a visual representation 3130. In this embodiment, there are multiple visual models 3120, each trained on a normalised part of a

vehicle and trained to determine if each part is classified as damaged or undamaged. The output visual representation 3130 comprises classifications of damaged or undamaged for each normalised part of the vehicle based on the images 3110. In some embodiments, the visual models 3120 can predict a rich representation (similar to or the same as that described in relation to other embodiments described herein), so for example can predict any or any combination of checks such as whether to repair or replace a part, the number of labour hours required, paint requirements, blend requirement, etc. and/or can also predict other visual features such as location of damage and/or type of damage. Any or any combination of these predictions can be relevant to determine whether and/or correlate with the vehicle was/being stationary or in movement when the damage occurred.

In parallel, the order sheet 3140 or other text/data containing information relating to how the damage to the vehicle occurred (or is claimed to have occurred) is input into a language model extractor 3150. The language model extractor 3150 is a natural language processing model trained to extract the circumstances of the damage 3160. The circumstances might include, for example, the point of impact, whether the vehicle was stationary or in motion, and the type of object with which the vehicle collided.

A classifier 3170 receives the output visual representation 3130 comprising classifications of damaged or undamaged for each normalised part of the vehicle and the circumstances of the damage 3160 and classifies whether the damage to each normalised part is consistent or not with the extracted circumstances 3160. For example, the location of damage may be inconsistent with the circumstances (e.g. it would be inconsistent to have damage on the front bumper and roof where the damage is not severe enough to indicate the vehicle rolled over) or the type of damage may be inconsistent (e.g. it would be inconsistent to have impact damage on a door and rust on a boot lid both being repaired due to an accident) or the severity of the damage may be inconsistent (e.g. it would be inconsistent to have scratches on the front bumper and a bigger dent on the bonnet, indicating the damage occurred at different times).

To train these models, the language model weights might be frozen while the visual model trains or both networks are trained jointly, end-to-end.

Referring now to Figure 32, an alternative embodiment to that of Figure 31 is presented in the situation where structure accident data is available, which will now be described in more detail.

As mentioned above, in this embodiment the accident data in the order sheet 3220 is stored (for example by the insurer computer system when an insurance claim is made) in a structured format that allows for direct processing of the structured data (i.e. no natural language processing is required to process the order sheet to extract the pertinent information regarding the circumstances of the damage).

The photos 3210 of the damaged vehicle and the order sheet 3220 including the details of the circumstances of the damage are provided to the model 3200. The circumstances of the damage to the vehicle are extracted from the structured data in an accident data extraction and injection step 3230. The circumstances might include, for example, the point of impact, whether the vehicle was stationary or in motion, and the type of object with which the vehicle collided.

The photos 3210 are processed by a visual damage classifier 3240, into which is also injected the circumstances of the damage to the vehicle, to predict whether there is damage to each part of the vehicle. A visual representation 3250 is output by the classifier 3240 including the classification of which parts are damaged or undamaged along with the circumstances of the damage. The representation can include the location of damage, the type of damage and the severity of damage.

The representation is provided to a further classifier 3260 which outputs a prediction of whether the damage to each part is consistent with the circumstances extracted from the order sheet 3220. For example, the location of damage may be inconsistent with the circumstances (e.g. it would be inconsistent to have damage on the front bumper and roof where the damage is not severe enough to indicate the vehicle rolled over) or the type of damage may be inconsistent (e.g. it would be inconsistent to have impact damage on a door and rust on a boot lid both being repaired due to an accident) or the severity of the damage may be inconsistent (e.g. it would be inconsistent to have scratches on the front bumper and a bigger dent on the bonnet, indicating the damage occurred at different times).

To train these models, the language model weights might be frozen while the visual model trains or both networks are trained jointly, end-to-end.

Referring now to Figure 33, there is shown a process 3300 for predicting damage to grouped auxiliary parts associated with a normalised panel according to an embodiment, which will now be described in more detail.

Auxiliary parts are parts of a vehicle such as any or any combination of: cameras, headlights, sensors, minor grilles, numberplates, emblems, fuel lids, brackets and such like. Some of these auxiliary parts tend to be difficult to determine accurately in images of a vehicle, and damage to these auxiliary parts can also tend to be difficult to determine from images of a vehicle.

Photos 3310 of the damage to the vehicle are input into the process 3300 and provided to an intermediary visual classifier 3320. The intermediary visual classifier 3320 comprises a plurality of per-part classifiers, each trained to determine a classification of damage to a normalised part of the vehicle for each of the input images 3310 and how severe the damage is. For example, the severity can range from 1 (scratches) to 5 (all areas affected).

In other embodiments, the intermediary visual classifier is trained to predict the location of the damage within the normalised part (for example: left/centre/right).

The output classifications of the intermediary visual classifier 3320 are provided to a multi-task learned classifier 3330, which is trained to simultaneously predict damage to multiple auxiliary parts of the vehicle per normalised part, based on the damage classifications to the normalised parts of the vehicle for the photos 3310.

As auxiliary parts, typically multiple auxiliary parts, can be associated with specific normalised parts (e.g. for a front bumper, auxiliary parts can include minor grilles, fog lamps, numberplates, emblems, parking sensors, impact bars, etc), a trained multi-task classifier can learn common features between all auxiliary parts on each normalised part/panel/region. This avoids training per-part models for auxiliary parts, which can be harder to train as not all vehicles have some auxiliary parts thus training data tends to be unbalanced or minimal.

The multi-task learned classifier 3330 can take the information determined by the visual classifier 3320 and use this to determine the likelihood of damage to one or more of the plurality of auxiliary parts associated with each normalised part based on for example the severity of the damage to each normalised part and whether the damage is present in a region of the normalised part in which the auxiliary part(s) are located.

In alternative embodiments, the visual classifier 3320 can perform a segmentation of the damage to each normalised part to identify more precisely where the damage on the vehicle is located, providing richer data for the multi-task learned classifier 3330 to predict whether damage has occurred to any auxiliary parts using this segmentation data. Further, in other embodiments, the visual classifier 3320 can perform a classification of the category of damage (e.g. scratch, dent, tear, rust, etc). to increase the richness of the data for the multi-task learned classifier 3330 to predict whether damage has occurred to any auxiliary parts.

Finally, an output classification is output 3340 by the multi-task learned classifier 3330 indicating whether any auxiliary parts are determined to be damaged. In some embodiments, the multi-task learned classifier 3330 takes just the output of the intermediary visual classifier 3320 as an input. In other embodiments, the multi-task learned classifier 3330 takes both the input images 3310 and the output of the intermediary visual classifier 3320 as inputs, for example the output of the intermediary visual classifier 3320 might comprise one or more segmentation masks of the images 3310, thus the images will be needed in addition to the segmentation masks to determine the classifications 3340).

In other embodiments, the intermediary visual classifier 3300 may not be required and just a multi-task classifier 3330 is used.

Referring now to Figure 34, an example multi-task architecture that can be used in the embodiment shown in Figure 33 is shown, and will now be described in more detail.

The multi-task architecture shown has both shared layers 3460, 3430, 3480 and task specific layers 3410, 3420, 3490, 3440, 3450, 3470. In this example, there are three tasks to be performed by the model, tasks A, B and C. Task A has some task-specific layers 3440, 3410 while similarly task B has some task-specific layers 3450, 3420 and task C has some task-specific layers 3470, 3490.

In the example embodiment of Figure 33, each of tasks A, B and C might relate to one of three auxiliary parts on one of the normalised parts.

Referring now to Figure 35, there is shown a process 3500 for jointly training across variably sized datasets, according to an embodiment, which will now be described in more detail.

In some cases, models 3530 are required to be used with different datasets 3510, 3520, for example for datasets from different geographies 3510, 3520. Further, where the different datasets 3510, 3520 are significantly different sizes relative to each other, or when one dataset is too small to train a model 3530 to output accurate results, it can prove hard to train a sufficiently accurate model 3530 for both datasets. One option is to train the model 3530 on the larger dataset 3520 and fine tune it on the smaller dataset 3510 in order to be able to output sufficiently accurate predictions for the smaller dataset 3510, but even this might fail if the smaller dataset 3510 is too dissimilar in size relative to the larger dataset 3520.

In this embodiment, a model 3530 is trained jointly on both datasets 3510, 3520 to evaluate the extent of the damage per normalised part, and then to predict a probability as to whether a panel should be replaced and the number of hours it would take to replace it. A multi-learning approach similar to that presented in other embodiments can be used to train the model 3530. The predictions can then be adapted for each dataset/geography to output final predictions 3530 per dataset 3510, 3520.

Further, in this embodiment, a domain confusion loss 3540 is injected into the model 3530 mid-way through the network to ensure that the network can't differentiate between data from the first dataset 3510 and the second dataset 3520, thus training the model 3530 to be generalised across both datasets 3510, 3520 (and therefore, for example, to work across both geographies). In alternative embodiments, no domain confusion loss 3540 is used.

Referring now to Figure 36, an alternative approach to working with multiple different sized datasets 3600 is presented according to another embodiment, which will now be described in more detail.

In this example, there are two datasets 3610, 3620 – the first geography dataset 3610 is much smaller than the second geography dataset 3620. If the most relevant data from the larger dataset 3620 can be identified, then it can be used to jointly train a model 3650 on the first dataset 3610 and a subset of the second dataset 3620. To identify the most relevant data from the larger dataset 3620, the most similar data points from the larger dataset 3620 are

selected (compared to the smaller dataset 3610) and used to populate a reduced second geography dataset 3640. In this embodiment, classifiers trained on the larger dataset 3620 are used to determine which data points are similar to the data in the smaller dataset 3610. This can be done by creating visual damage vectors for all the data in both the larger and smaller datasets 3610, 3620 and then finding the data points in the vectors/data points created for the larger dataset 3620 that are most similar/close to the vectors/data points created for the smaller dataset 3610. To determine how “close” the vectors/data points are, in terms of visual damage vectors, a simple distance measure can be used (for example Euclidian, Mean Squared, etc.) or a simple classifier can be trained (for example using linear regression or a support vector machine) to separate the two datasets and determine which points in the larger dataset don’t get separated from the smaller dataset. If even a small amount of the larger dataset 3620 is selected 3640 to augment the first dataset 3610, then it is possible that a substantially improved model 3660 can be generated following joint training across the first dataset 3610 and second dataset 3640.

Referring to Figure 37, there is shown an example domain confusion loss architecture 3700 that can be used with other embodiments, which will now be described in more detail.

Two networks 3720, 3790 are shown, one network 3730 for labelled images 3760 and one network 3790 for unlabelled images 3780. Each network has a number of convolutional layers 3730, 3710, 3792, 3791 and fully connected layers 3714, 3713, 3712, 3711, 3795, 3794, 3793, 3792, 3770 and have a domain loss 3740 injected in between layers 3770, 3794, 3714, 3712. Additionally, there is a classification loss 3750 injected into one of the networks 3720. Additionally, there is an adaption layer 3713, 3795 in each network 3720, 3790. In this example, the two domains being considered are labelled and unlabelled images, and the goal of the domain confusion loss 3740 is to make sure that at the fc8 layer 3714, 3770 the network can’t differentiate between data that came from the labelled domain 3760 and data that came from the unlabelled domain 3780. It does this by having a small classifier on top of the fc8 layer 3714, 3770 which tried to predict whether a particular image came from the labelled or unlabelled domains. If the small classifier predicts correctly, it penalises the fc8 layer 3714, 3770 in the network. In this way, while training, the network will learn that in the fc8 layer 3714, 3770 it should not be possible to differentiate data from the two domains 3760, 3780, thus leading to better generalisation across the two domains 3760, 3780.

It should be understood by the reader that the cropping and/or segmentation techniques described in some embodiments herein can be adapted for use in other embodiments and/or aspects described herein, for example in the determination of whether a part is damaged or undamaged and/or the determination of damage to auxiliary parts. It follows that these techniques are thereby disclosed for use in relation with any aspects/embodiments herein.

It should also be understood by the reader that the multi-image learning approaches described in embodiments and/or aspects herein can be incorporated into other aspects/embodiment to substantially improve accuracy for all classification or models across multiple images and/or other input data. It follows that these approaches are thereby disclosed
5 for use in relation with any aspects/embodiments herein.

The above described embodiments will be able to be implemented using a variety of software and hardware, all within the intended scope of the aspects and embodiments set out herein. For example, the software may be written and deployed in a variety of ways using a variety of languages and implementations. The hardware that may be used to deploy the
10 software and methods/systems described herein according to any embodiment and/or aspect may be a variety of different types and configurations of hardware and hardware architecture including the use of any or any combination of: handheld devices, mobile 'phones, smart 'phones, cameras, imaging devices, tablet computers, desktop computers; laptop computers; virtual desktops; servers; distributed servers; distributed computing arrangements; and cloud
15 systems.

By models and neural networks, the use of machine learning and artificial intelligence techniques are intended to be referred to, including the use of deep learning techniques and convolutional and other neural networks.

Machine learning is the field of study where a computer or computers learn to perform
20 classes of tasks using the feedback generated from the experience or data gathered that the machine learning process acquires during computer performance of those tasks.

Typically, machine learning can be broadly classed as supervised and unsupervised approaches, although there are particular approaches such as reinforcement learning and semi-supervised learning which have special rules, techniques and/or approaches.
25 Supervised machine learning is concerned with a computer learning one or more rules or functions to map between example inputs and desired outputs as predetermined by an operator or programmer, usually where a data set containing the inputs is labelled.

Unsupervised learning is concerned with determining a structure for input data, for example when performing pattern recognition, and typically uses unlabelled data sets.
30 Reinforcement learning is concerned with enabling a computer or computers to interact with a dynamic environment, for example when playing a game or driving a vehicle.

Various hybrids of these categories are possible, such as "semi-supervised" machine learning where a training data set has only been partially labelled. For unsupervised machine learning, there is a range of possible applications such as, for example, the application of
35 computer vision techniques to image processing or video enhancement. Unsupervised machine learning is typically applied to solve problems where an unknown data structure might be present in the data. As the data is unlabelled, the machine learning process is required to

operate to identify implicit relationships between the data for example by deriving a clustering metric based on internally derived information. For example, an unsupervised learning technique can be used to reduce the dimensionality of a data set and attempt to identify and model relationships between clusters in the data set, and can for example generate measures of cluster membership or identify hubs or nodes in or between clusters (for example using a technique referred to as weighted correlation network analysis, which can be applied to high-dimensional data sets, or using k-means clustering to cluster data by a measure of the Euclidean distance between each datum).

Semi-supervised learning is typically applied to solve problems where there is a partially labelled data set, for example where only a subset of the data is labelled. Semi-supervised machine learning makes use of externally provided labels and objective functions as well as any implicit data relationships. When initially configuring a machine learning system, particularly when using a supervised machine learning approach, the machine learning algorithm can be provided with some training data or a set of training examples, in which each example is typically a pair of an input signal/vector and a desired output value, label (or classification) or signal. The machine learning algorithm analyses the training data and produces a generalised function that can be used with unseen data sets to produce desired output values or signals for the unseen input vectors/signals. The user needs to decide what type of data is to be used as the training data, and to prepare a representative real-world set of data. The user must however take care to ensure that the training data contains enough information to accurately predict desired output values without providing too many features (which can result in too many dimensions being considered by the machine learning process during training and could also mean that the machine learning process does not converge to good solutions for all or specific examples). The user must also determine the desired structure of the learned or generalised function, for example whether to use support vector machines or decision trees.

The use of unsupervised or semi-supervised machine learning approaches are sometimes used when labelled data is not readily available, or where the system generates new labelled data from unknown data given some initial seed labels.

Machine learning may be performed through the use of one or more of: a non-linear hierarchical algorithm; neural network; convolutional neural network; recurrent neural network; long short-term memory network; multi-dimensional convolutional network; a memory network; fully convolutional network or a gated recurrent network allows a flexible approach when generating the predicted block of visual data. The use of an algorithm with a memory unit such as a long short-term memory network (LSTM), a memory network or a gated recurrent network can keep the state of the predicted blocks from motion compensation processes performed on the same original input frame. The use of these networks can improve

computational efficiency and also improve temporal consistency in the motion compensation process across a number of frames, as the algorithm maintains some sort of state or memory of the changes in motion. This can additionally result in a reduction of error rates.

Developing a machine learning system typically consists of two stages: (1) training and
 5 (2) production. During training, the parameters of the machine learning model are iteratively changed to optimise a particular learning objective, known as the objective function or the loss. Once the model is trained, it can be used in production, where the model takes in an input and produces an output using the trained parameters.

During the training stage of neural networks, verified inputs are provided, and hence it
 10 is possible to compare the neural network's calculated output to then the correct the network is need be. An error term or loss function for each node in neural network can be established, and the weights adjusted, so that future outputs are closer to an expected result. Backpropagation techniques can also be used in the training schedule for the or each neural network.

15 The model can be trained using backpropagation and forward pass through the network. The loss function is an objective that can be minimised, it is a measurement between the target value and the model's output.

The cross-entropy loss may be used. The cross-entropy loss is defined as

$$20 \quad L_{CE} = - \sum_{c=1}^C y * \log(s)$$

where C is the number of classes, $y \in \{0,1\}$ is the binary indicator for class c , and s is the score for class c .

In the multitask learning setting, the loss will consist of multiple parts. A loss term for
 25 each task.

$$L(x) = \lambda_1 L_1 + \lambda_2 L_2$$

where L_1, L_2 are the loss terms for two different tasks and λ_1, λ_2 are weighting terms.

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

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

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

Claims

1. A computer-implemented method for determining a measure of consistency between one or more damage indications to one or more parts of a vehicle, comprising the steps of:

receiving a plurality of images of the vehicle;

determining, per part, one or more classifications for each of the plurality of images using one or more trained models, wherein each classification comprises at least one indication of damage to at least one part;

determining a measure of consistency between the one or more indications of damage; and

outputting the measure of consistency of each of the parts of the vehicle wherein the indication of damage to at least two of the parts is determined to be correlated.

2. The method of claim 1, wherein determining whether the at least one indication of damage to at least one part has a measure of consistency above a predetermined threshold that indicates the at least one indication of damage was caused in a single accident and/or incident.

3. The method of claim 2, wherein determination of a single accident and/or incident comprises any one or any combination of: location of the one or more indications of damage on the vehicle; type of damage; severity of damage; and/or one or more additional structured queries comprising point and/or of impact; vehicle motion state and/or manually generated text data.

4. The method of any preceding claim further comprising receiving accident data, wherein the accident data comprises data about the cause of the damage to the vehicle.

5. The method of claim 4 further comprising:

processing the accident data to predict damage to one or more parts of the vehicle; and

determining a correlation between the determined at least one indications of damage and the predicted damage.

6. The method of any preceding claim, wherein the output further comprises one or more indications of damage and/or relevant images corresponding to one or more vehicle parts determined to have a measure of consistency above a predetermined threshold.

7. The method of claims 5 to 6, wherein processing the accident data comprises the use of any or any combination of: Natural Language Processing; a neural network; a convolutional neural network; a recurrent neural network; and/or one or more trained classifiers.

8. The method of any preceding claim further comprising determining a plurality of parts of the vehicle that are represented in the plurality of images of the vehicle comprising the use of a plurality of classifiers.

9. The method of claim 8, wherein each one of the plurality of classifiers is operable to detect each of the parts of the vehicle.

10. The method of any preceding claim, further comprising determining one or more sets of relevant images of each of the plurality of parts from the plurality of images of the vehicle.

11. The method of claim 10, further comprising determining a damage state of each of the parts of the vehicle using the one or more relevant images of each of the plurality of parts of the vehicle, the one or more damage states being determined as one or more quantitative values.

12. The method of any preceding claim, wherein determining the one or more classifications and/or determining the one or more damage states comprises using a multi-instance learned approach.

5 13. The method of any preceding claim, wherein determining the one or more classifications and/or determining the one or more damage states comprises using a graph neural network.

10 14. The method of any preceding claim, wherein determining the one or more classifications for each of the plurality of images further comprises determining one or more feature representations

15 15. A computer-implemented method of training a neural network for determining a measure of consistency between one or more damage indications to one or more parts of a vehicle, including neurons, each neuron being associated with weights, wherein the method comprises:

obtaining training inputs, the training inputs comprising images of vehicles, classifications of damage to vehicles and/or a damage consistency threshold;

20 for each training input, selecting one or more neurons based on their respective probability;

adjusting the weights of the selected neurons such that the selected neurons are substantially operable to identify one or more indications of damage;

processing the training input with the neural network to generate a predicted output;

25 adjusting the weights based on the predicted output.

16. The method of claims 15, wherein a plurality of neural networks are trained for each of the parts of a vehicle.

Traditional Insurance Claim Process

100

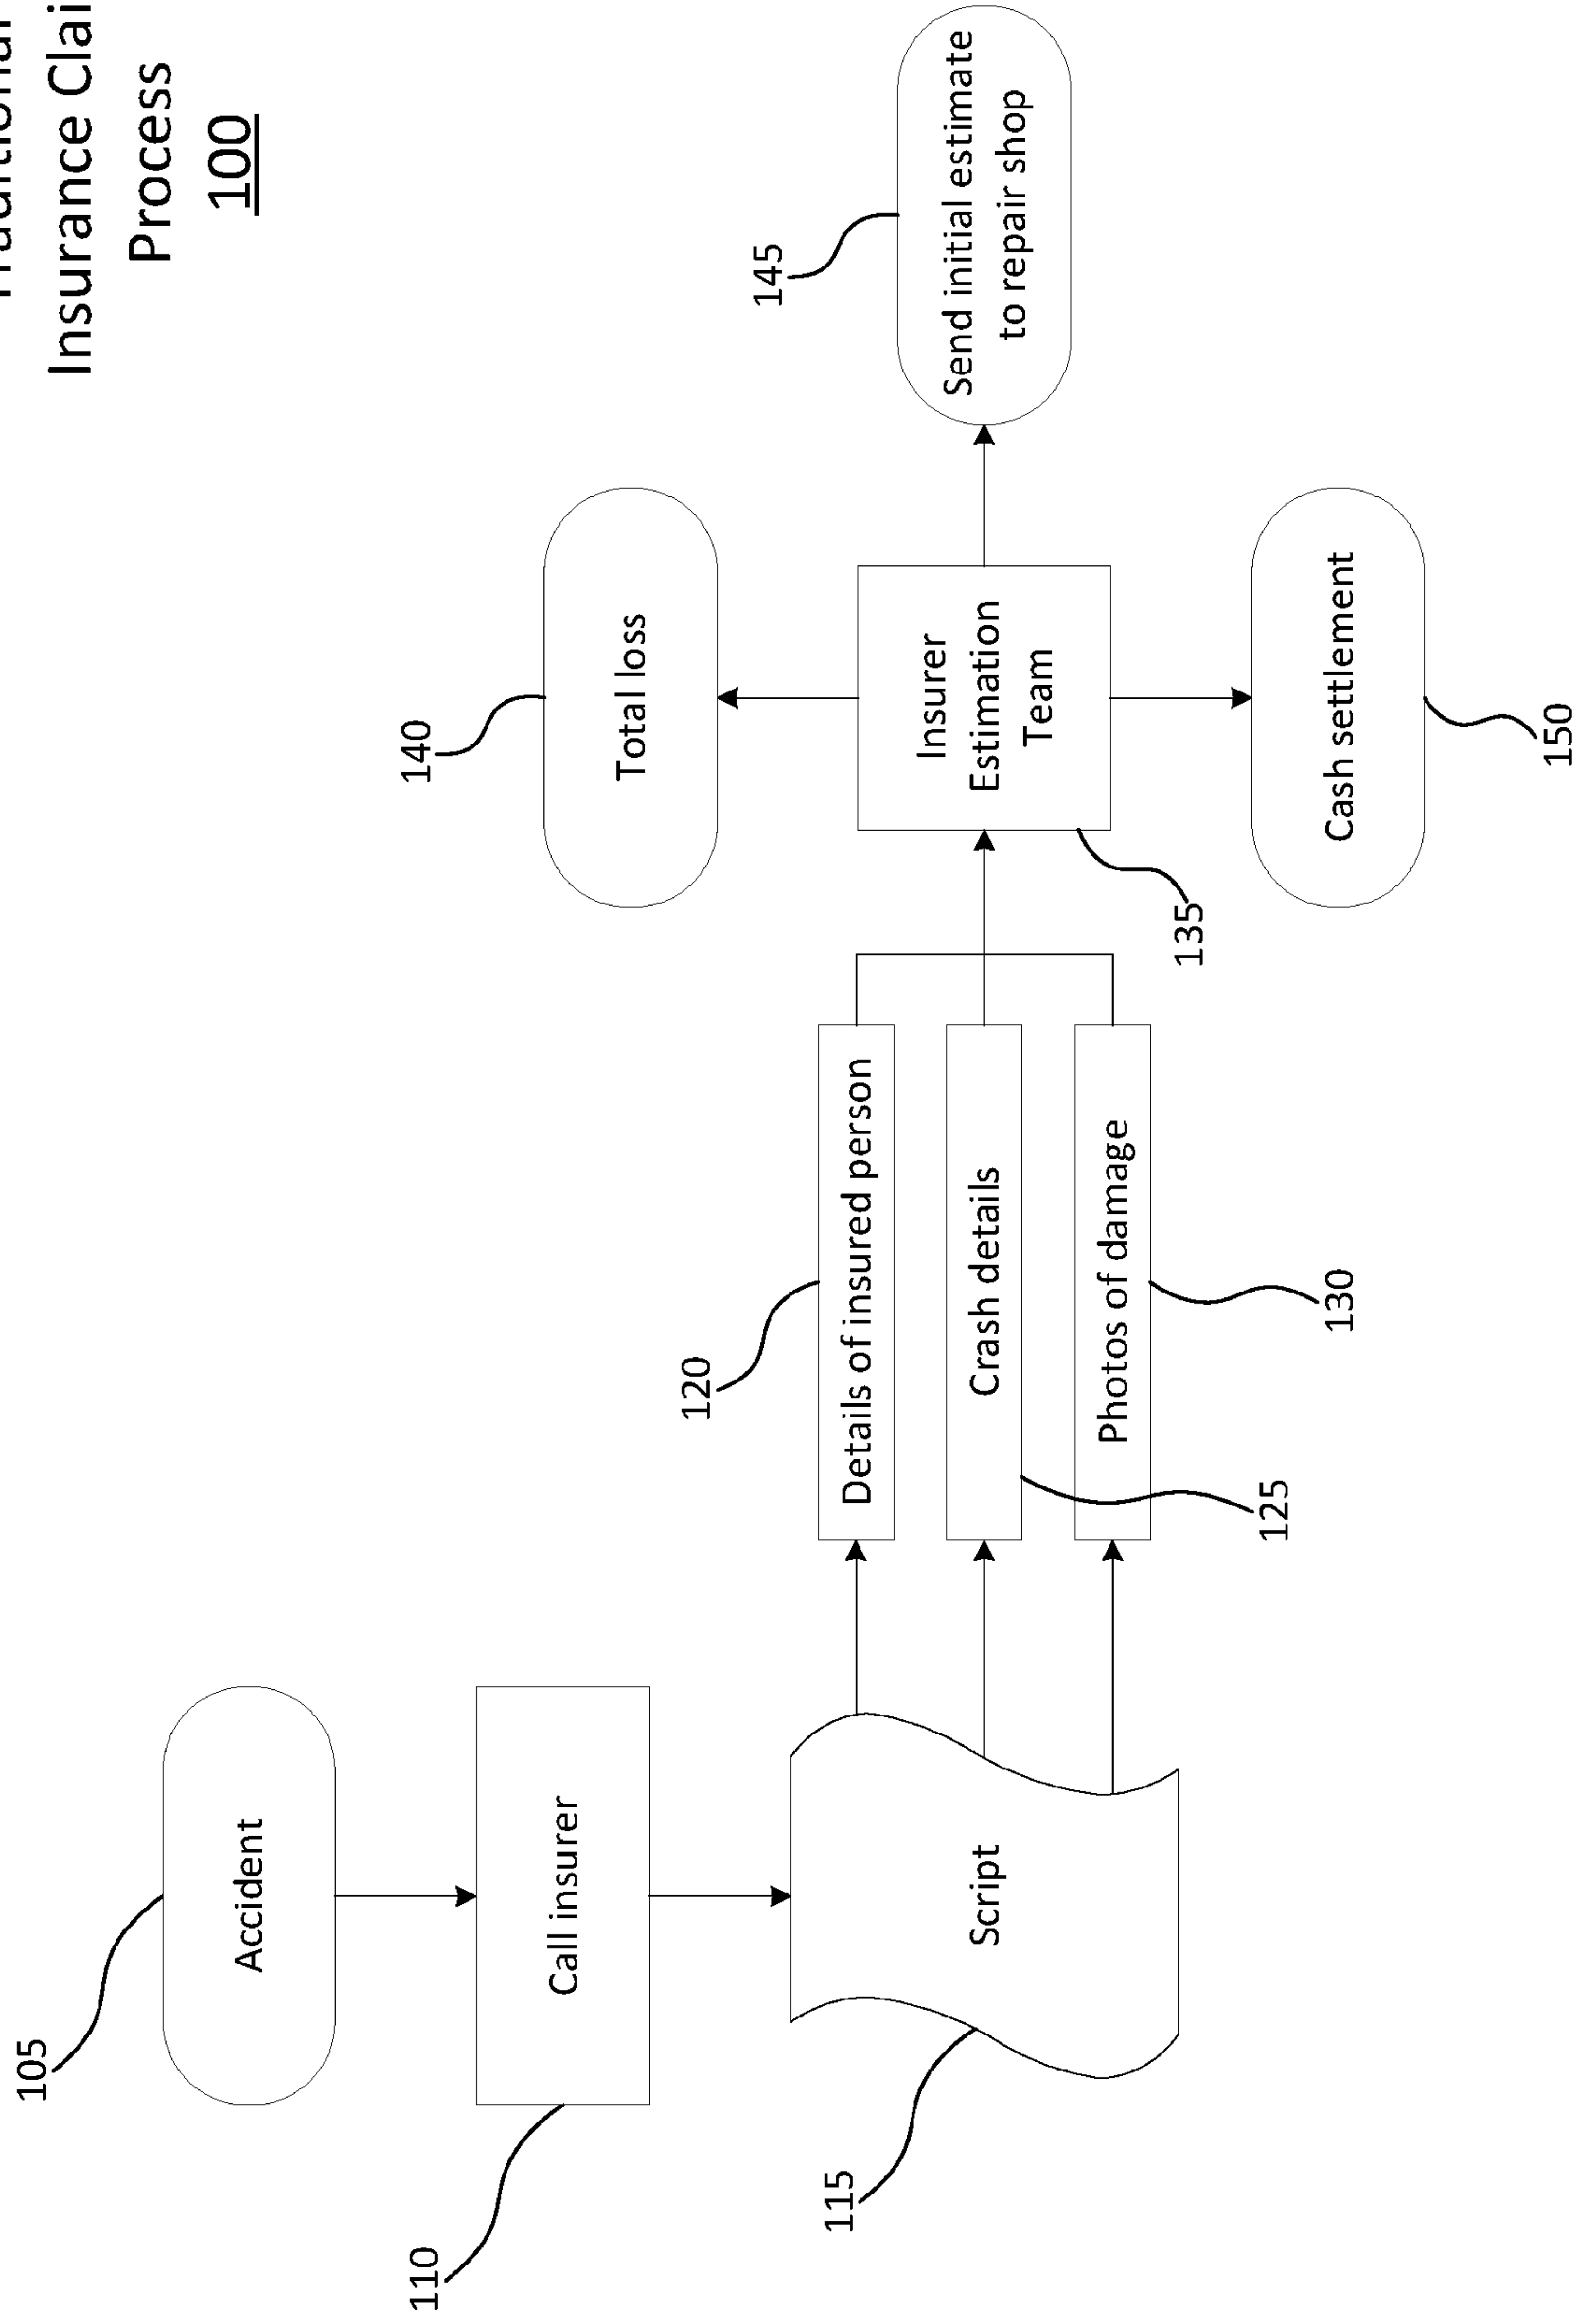


Figure 1

Near Real-Time Estimation Process

200


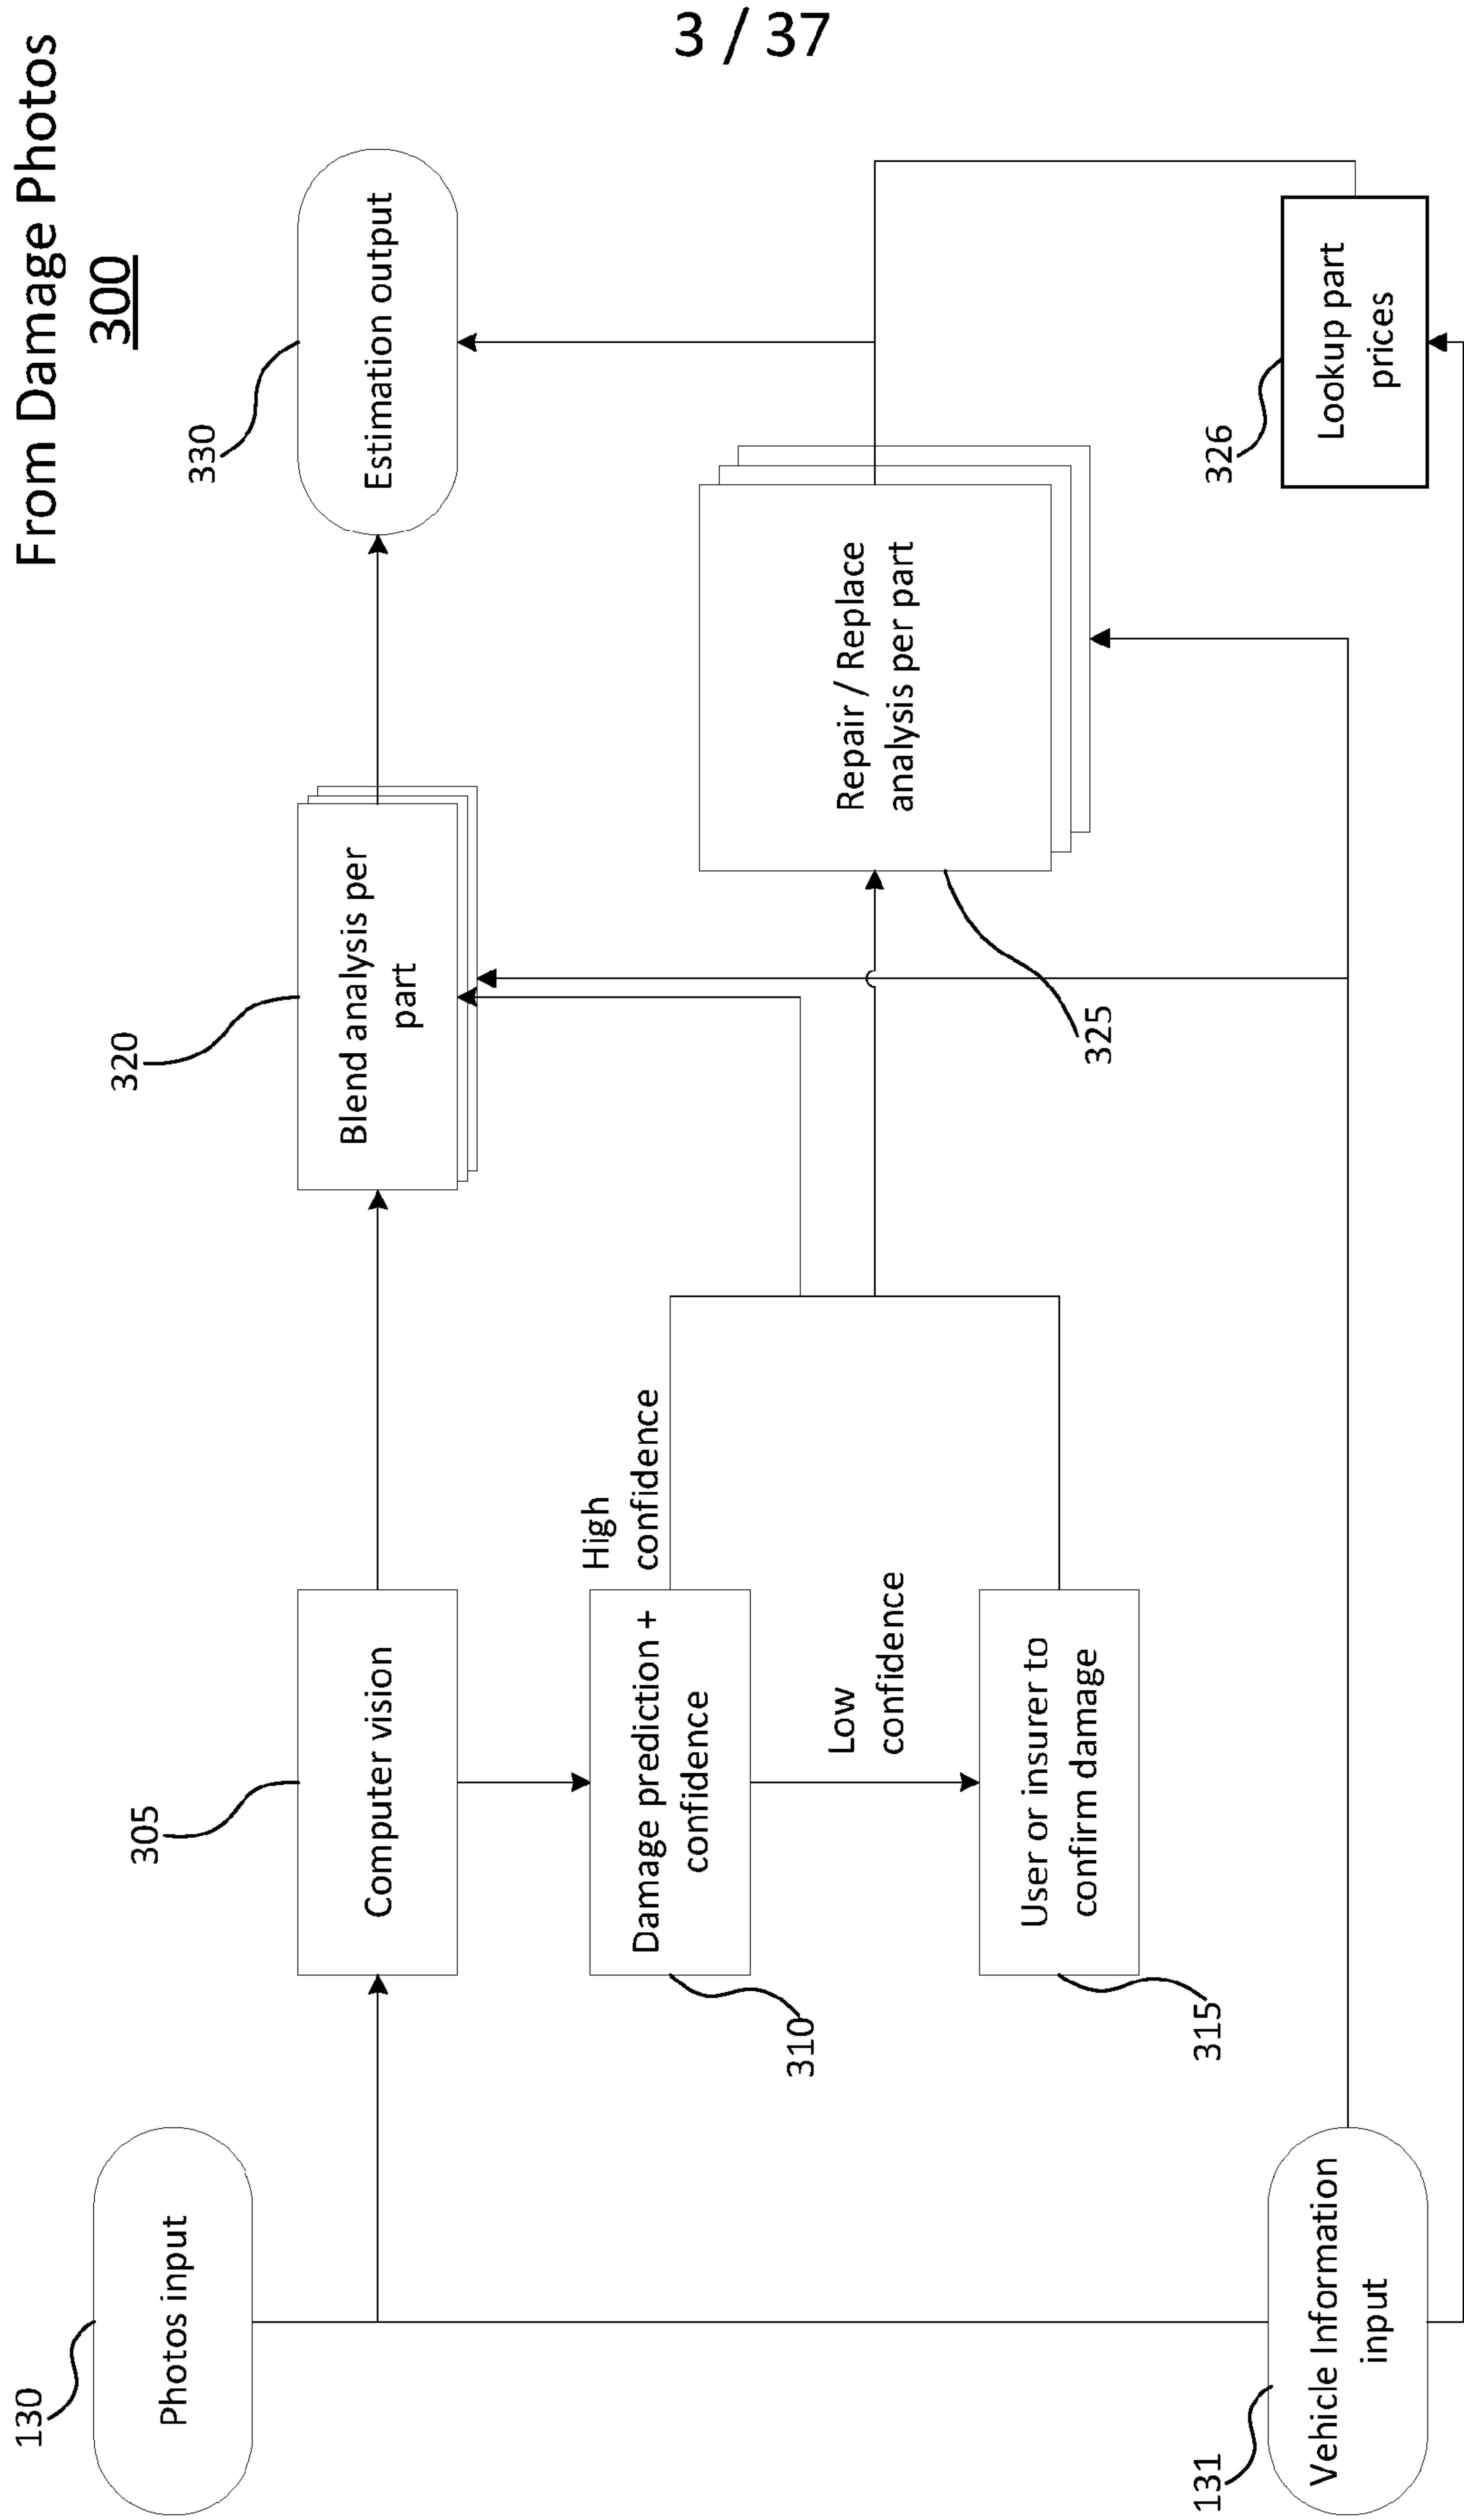


Figure 2

Determining Estimate From Damage Photos

300



3 / 37

Figure 3

Predicting Total Loss
400

4 / 37

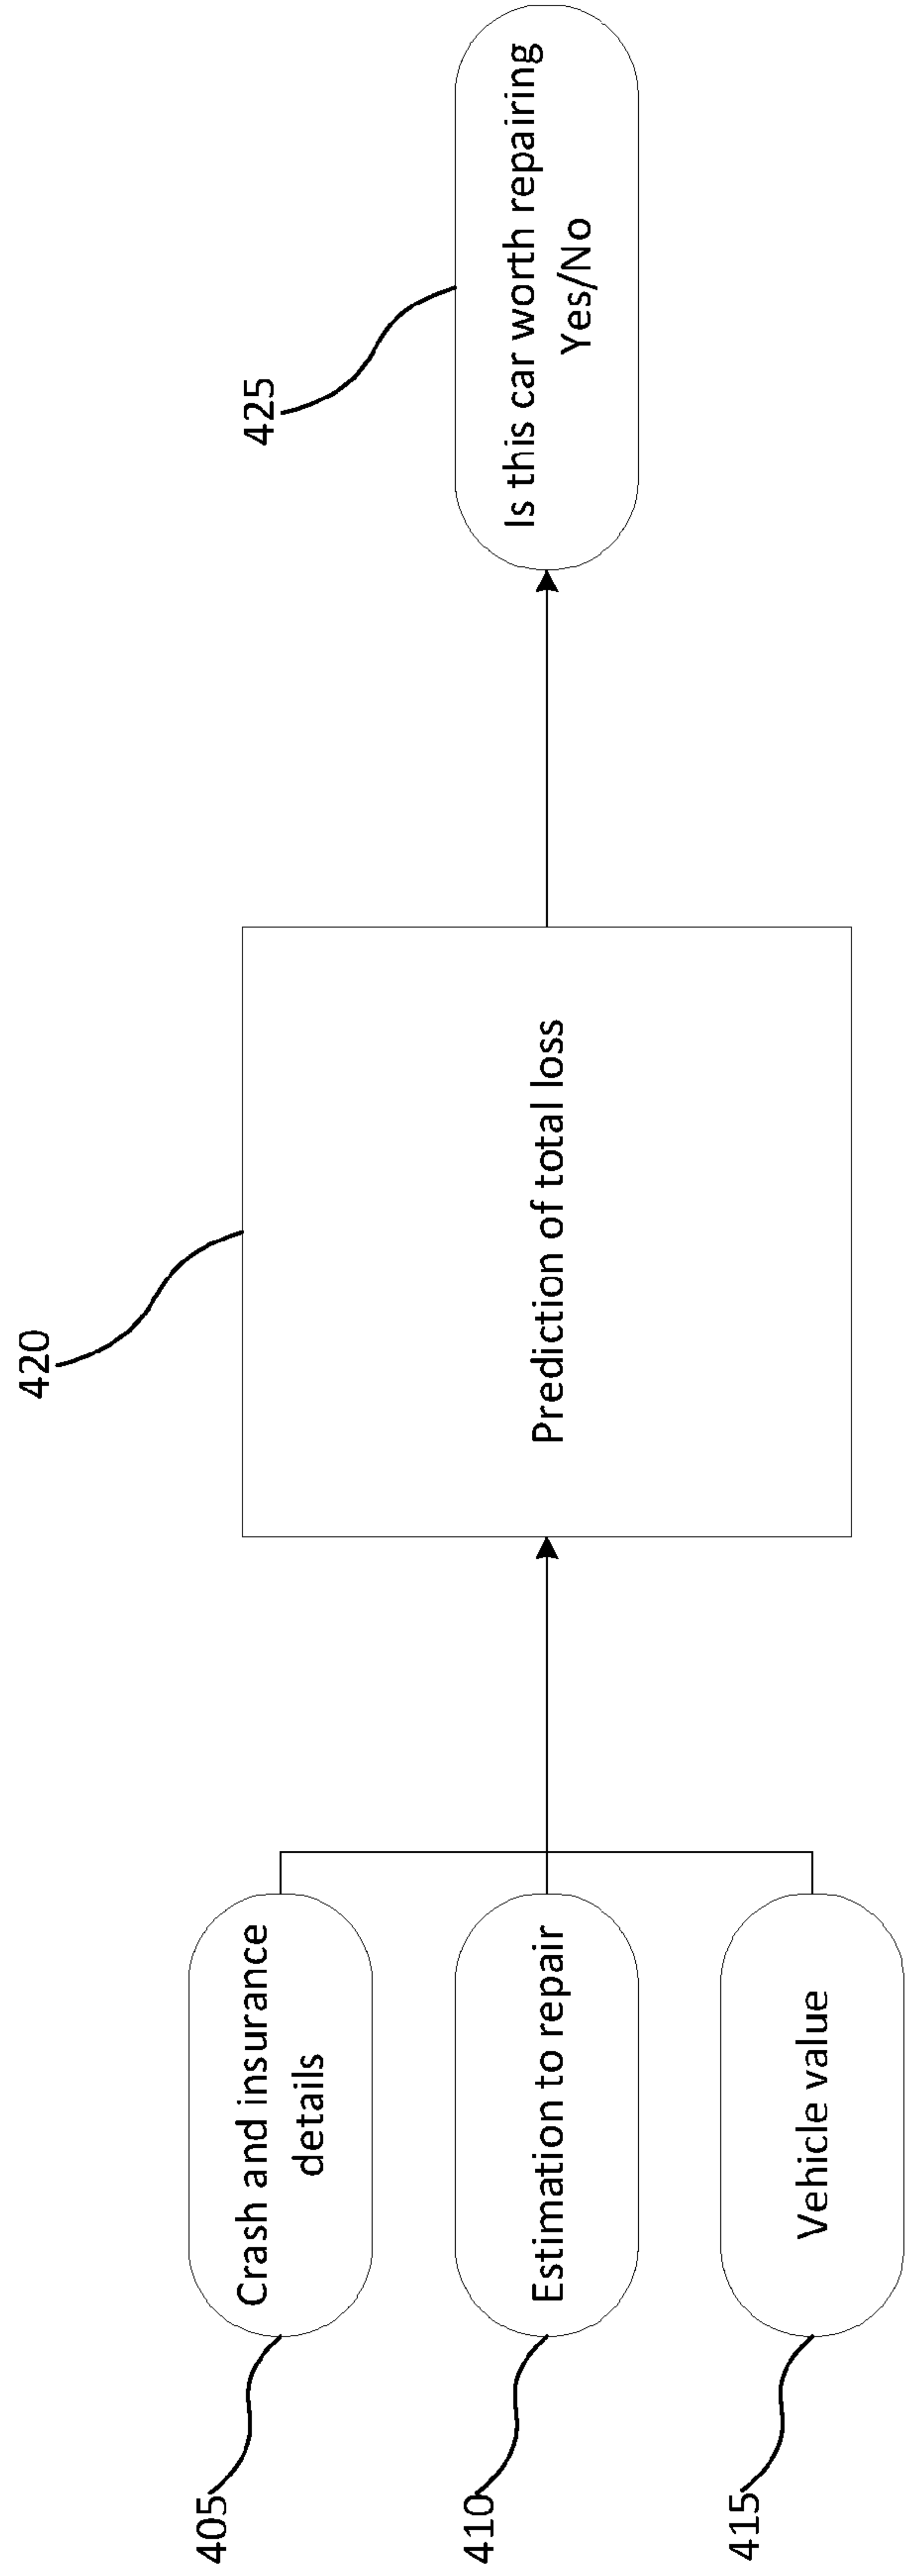


Figure 4

Estimating Repair Cost From Photos

500

5 / 37

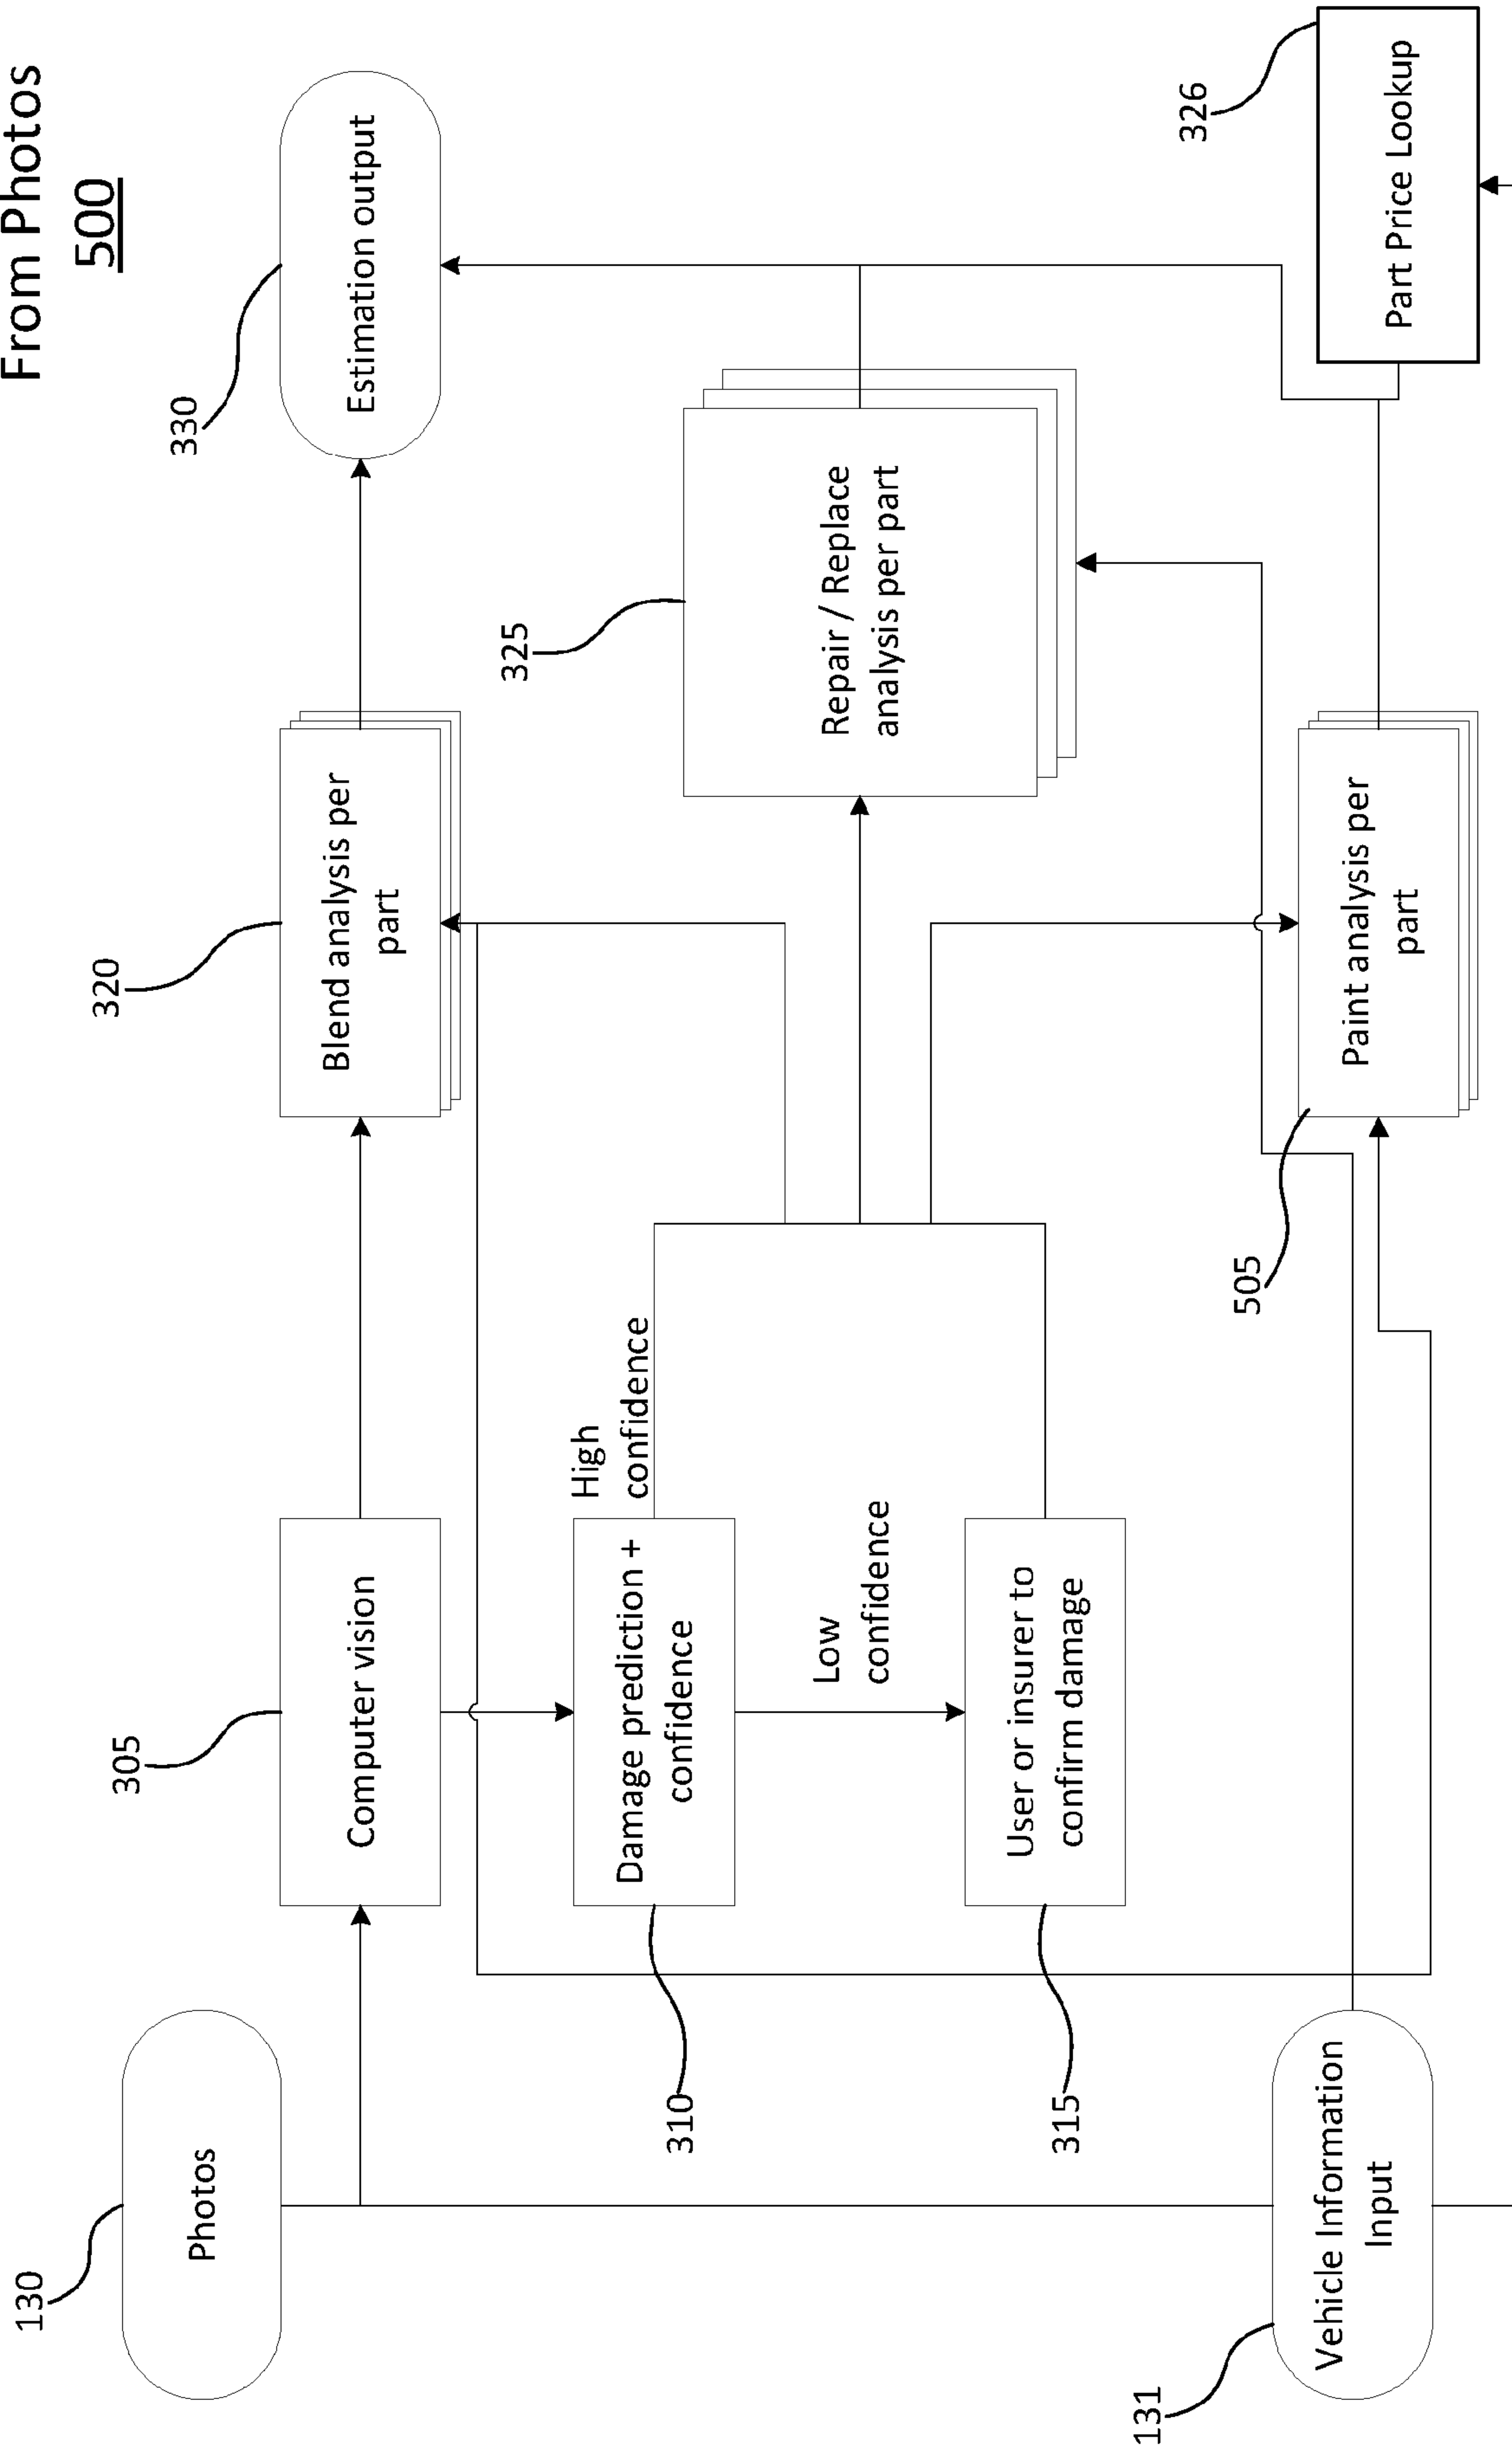


Figure 5

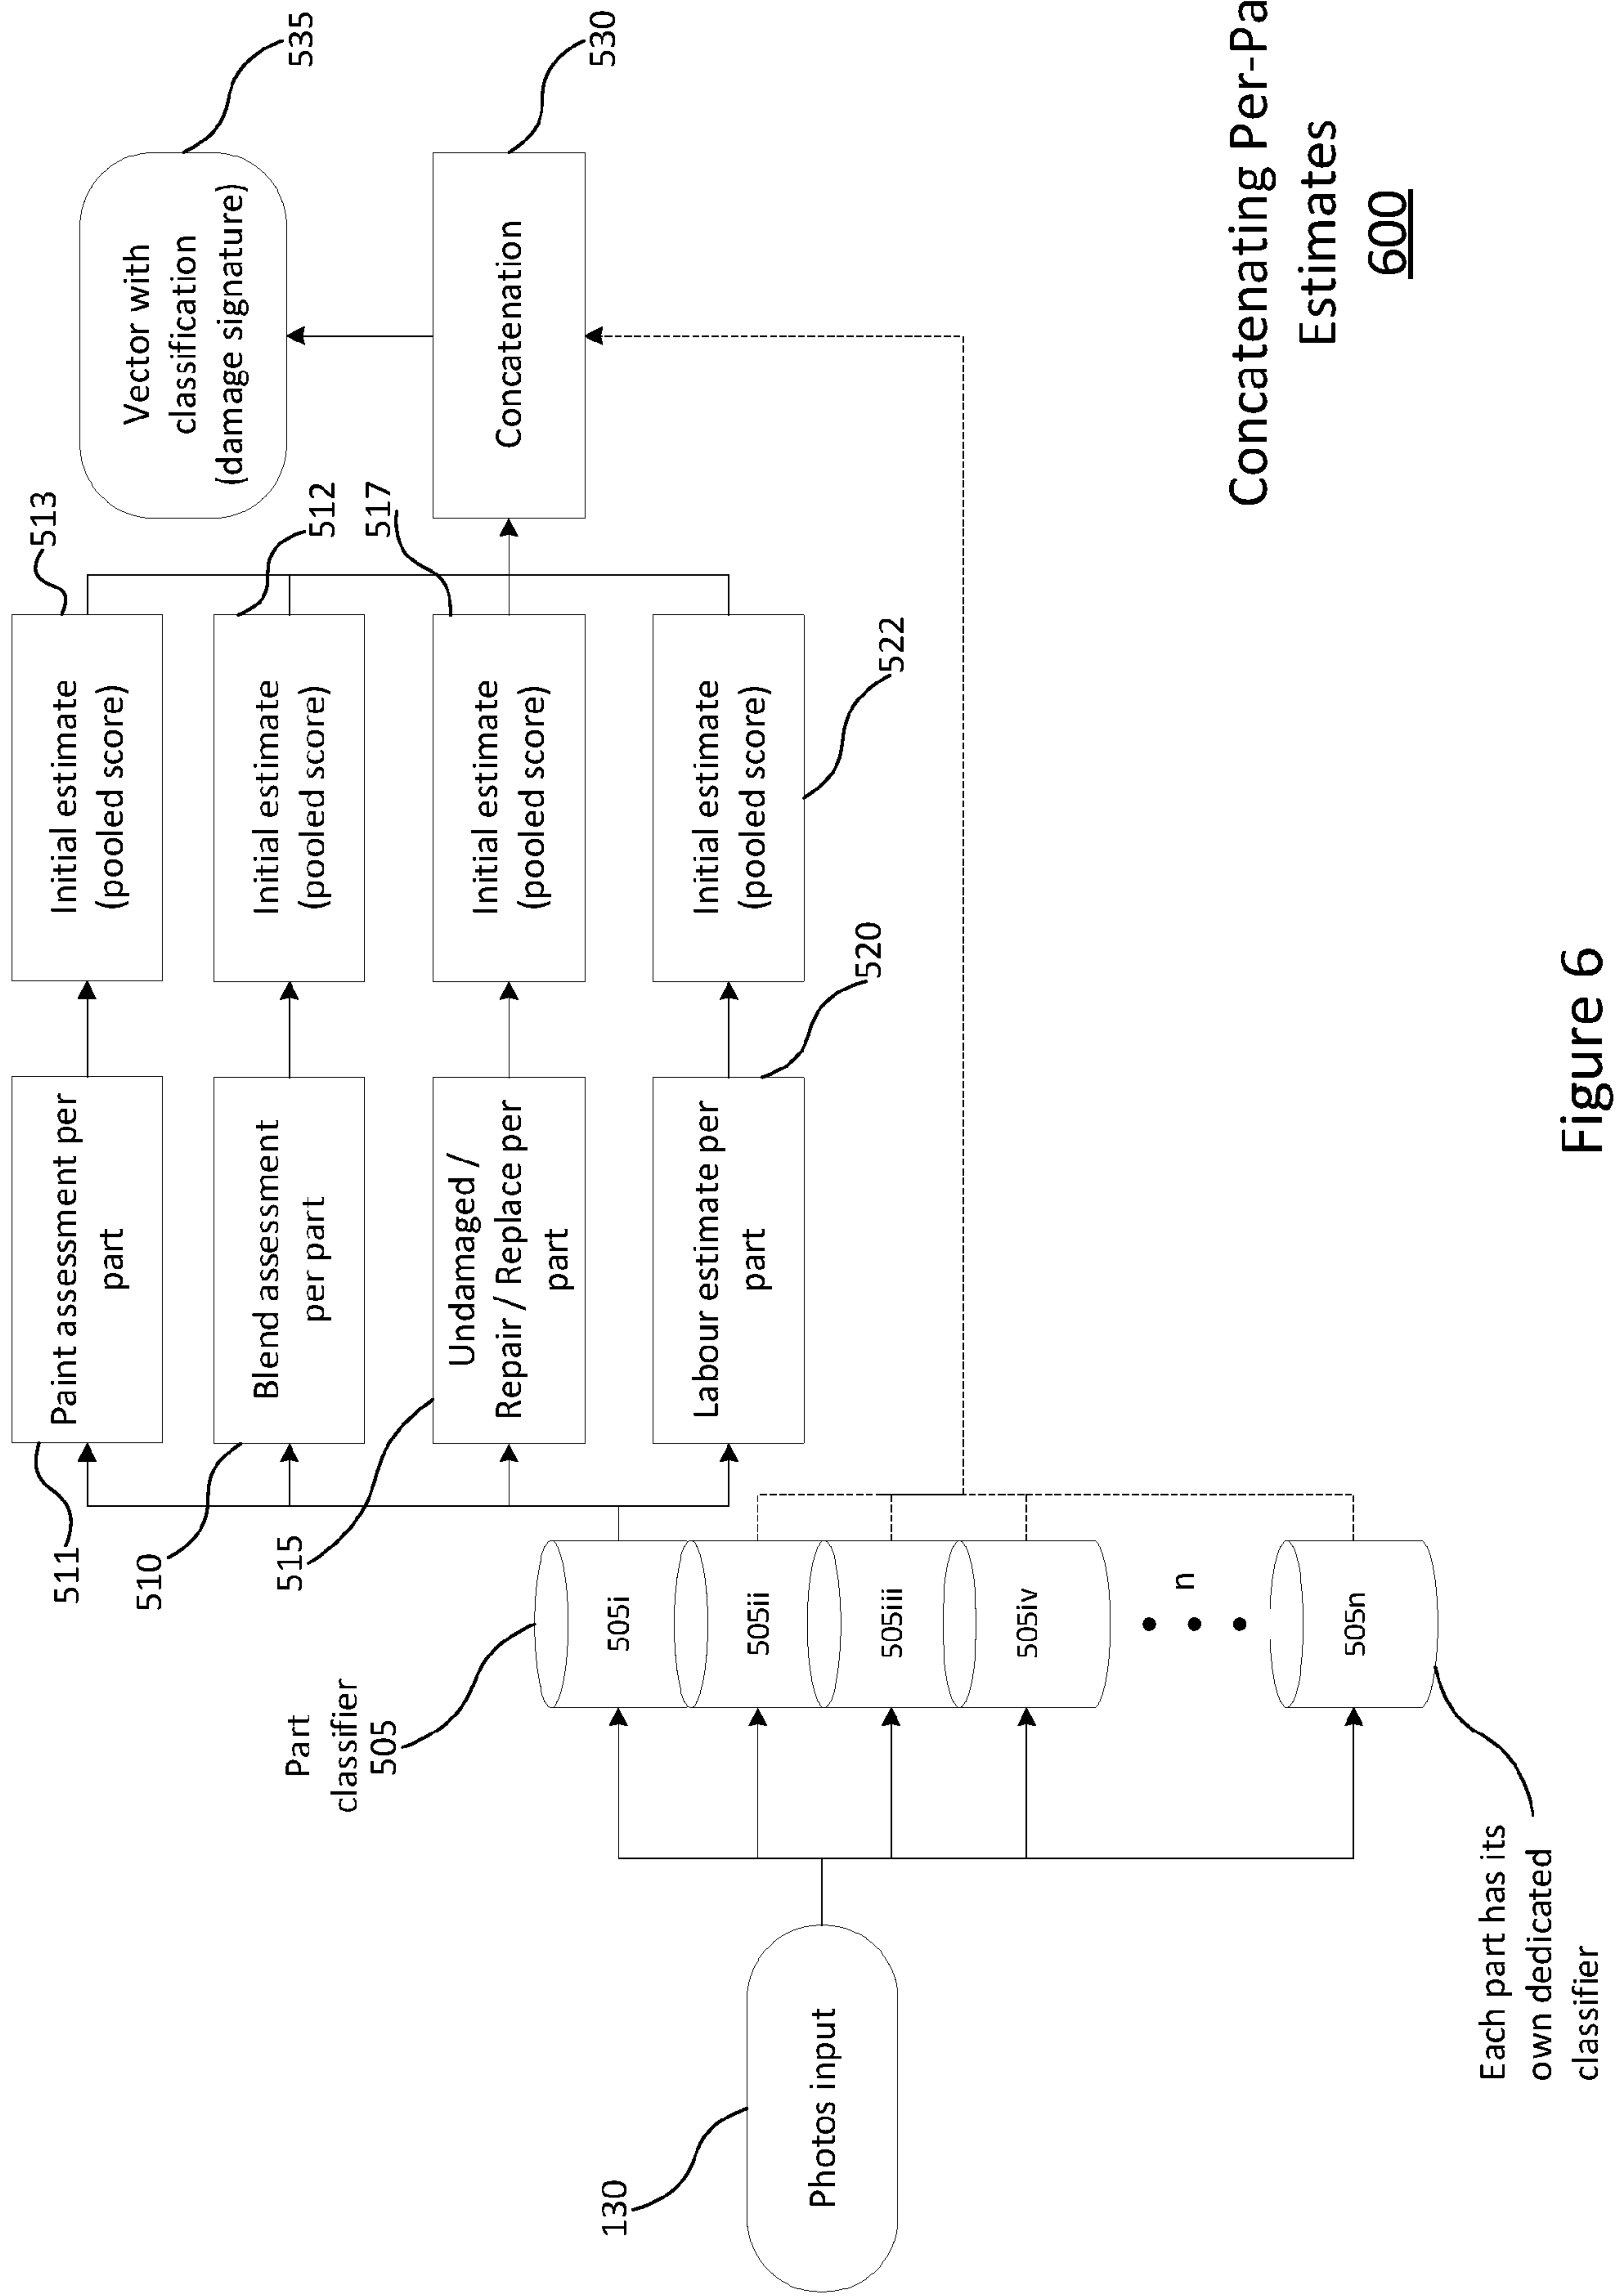


Figure 6

Determining Paint
Blending Requirement
700




Figure 7

Determining Whether Part is
Undamaged or Repairable or
Needs Replacing
800


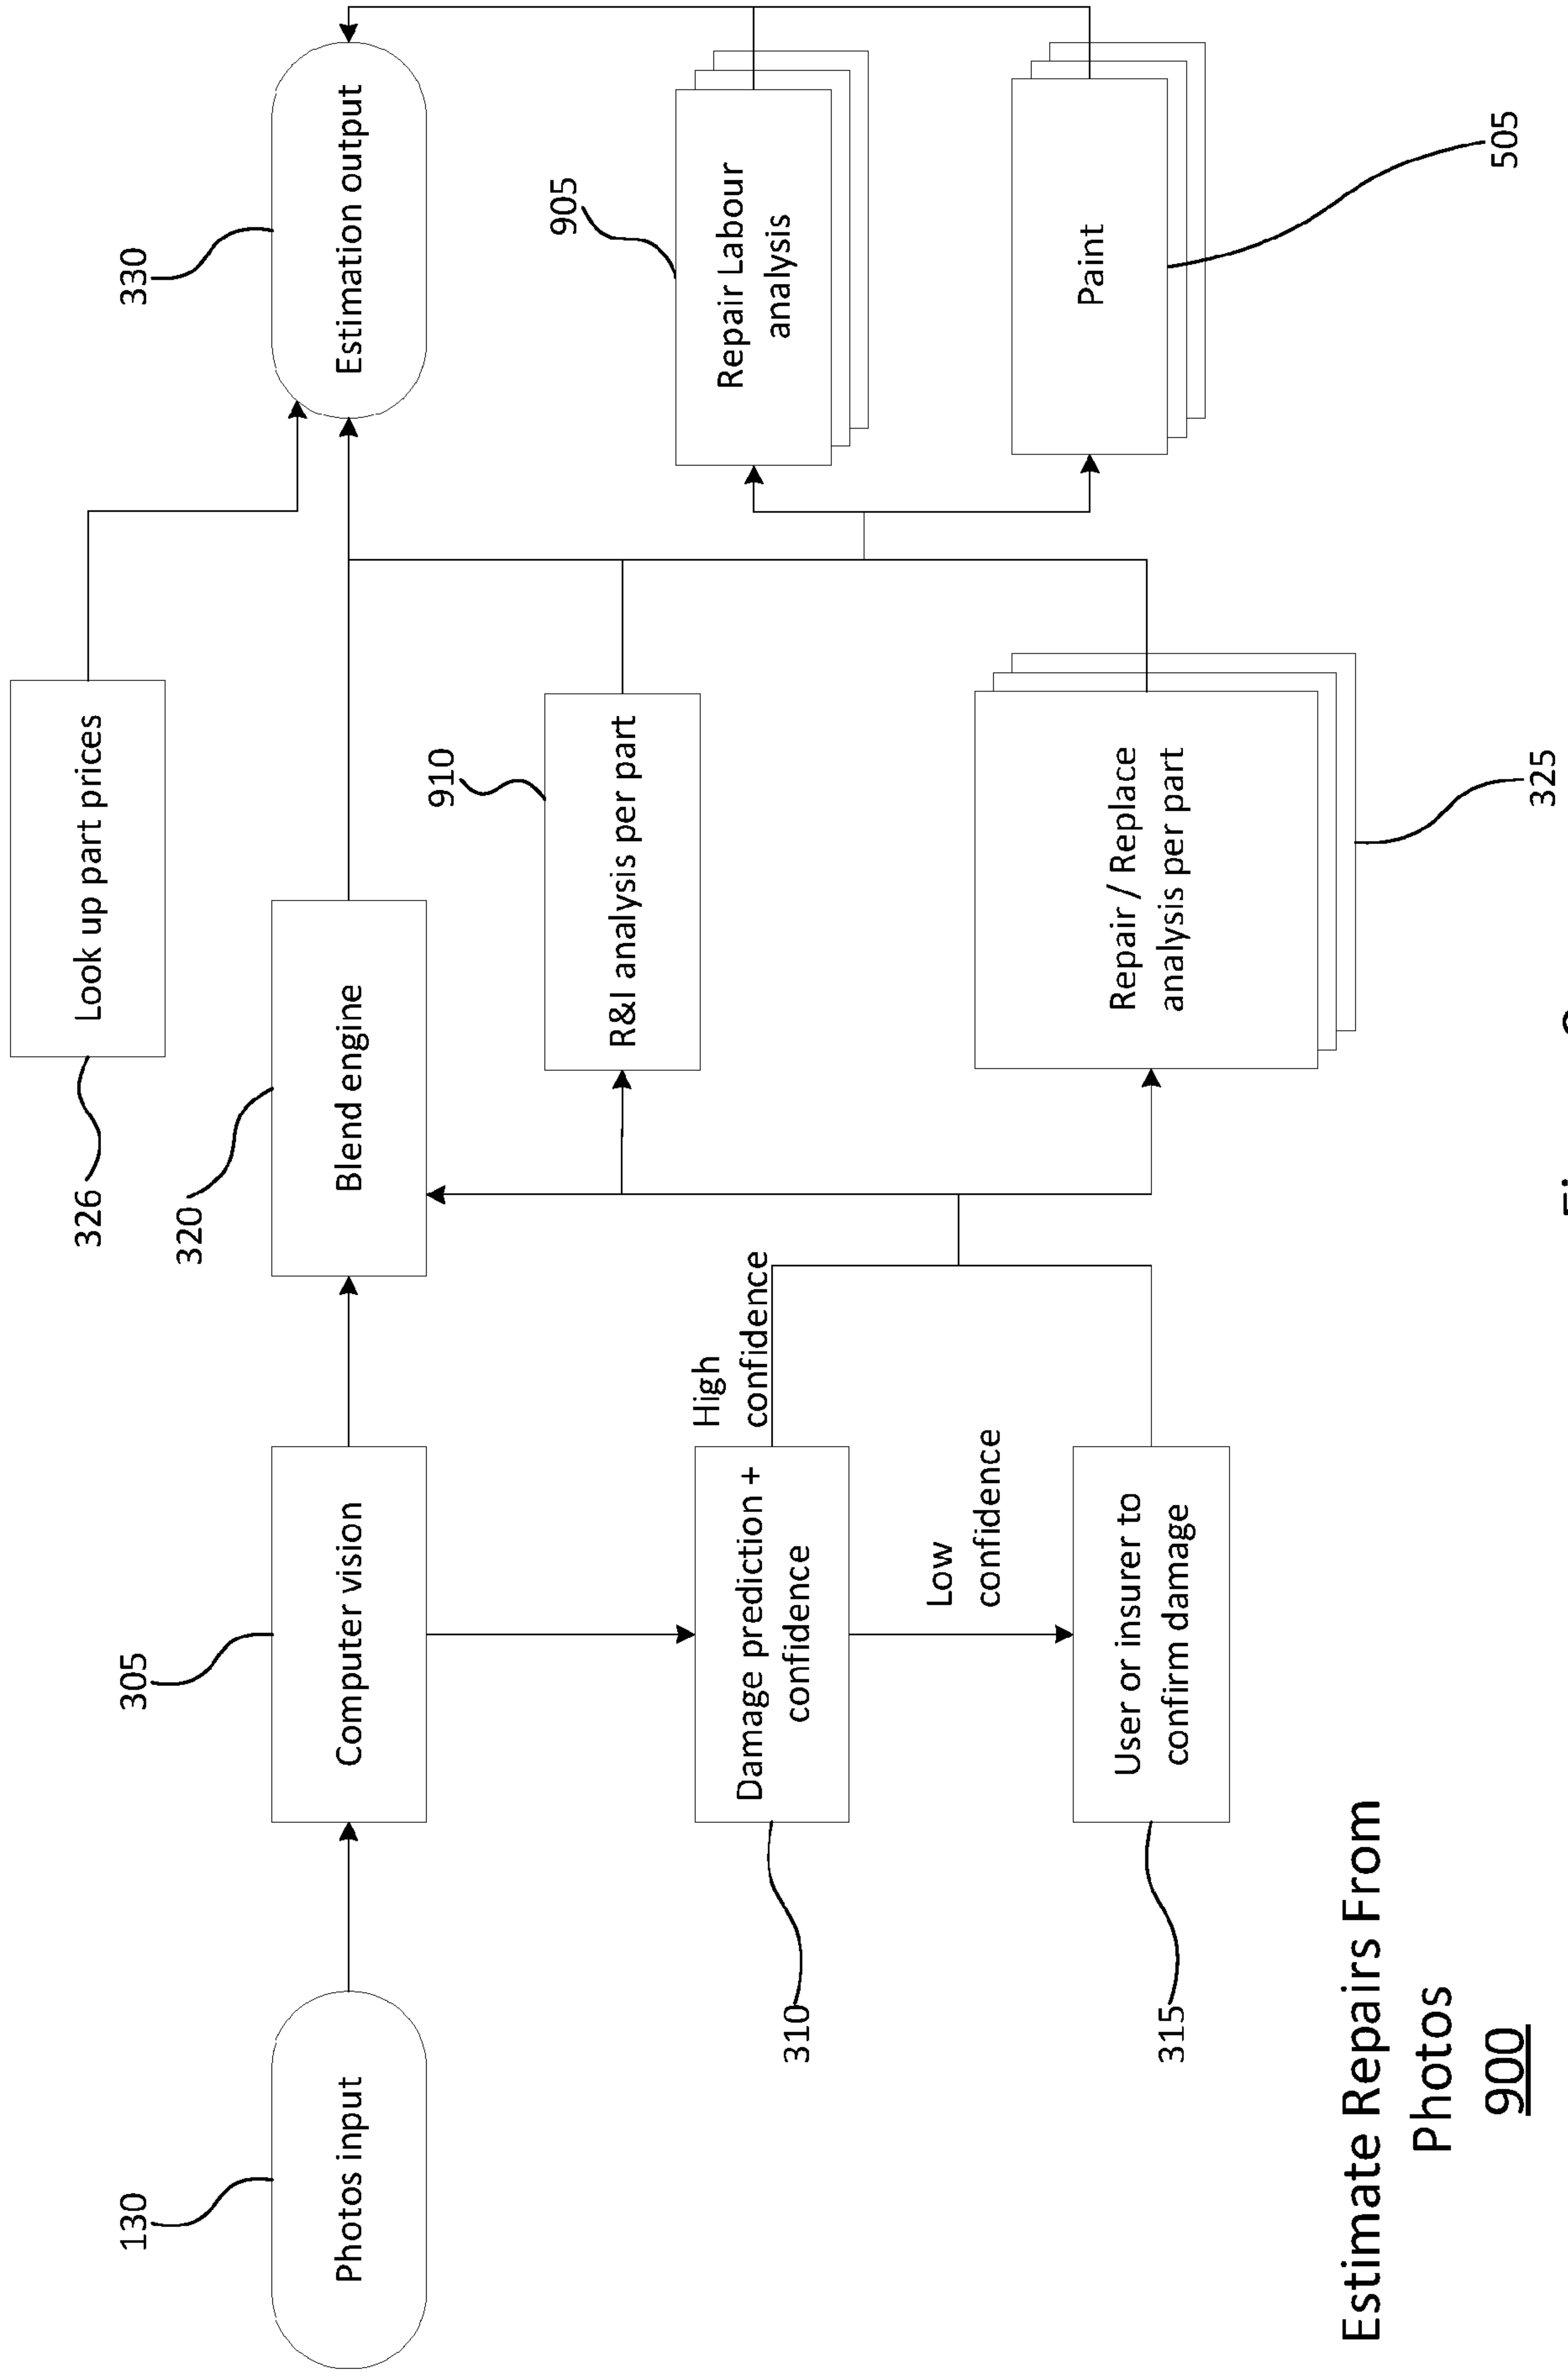
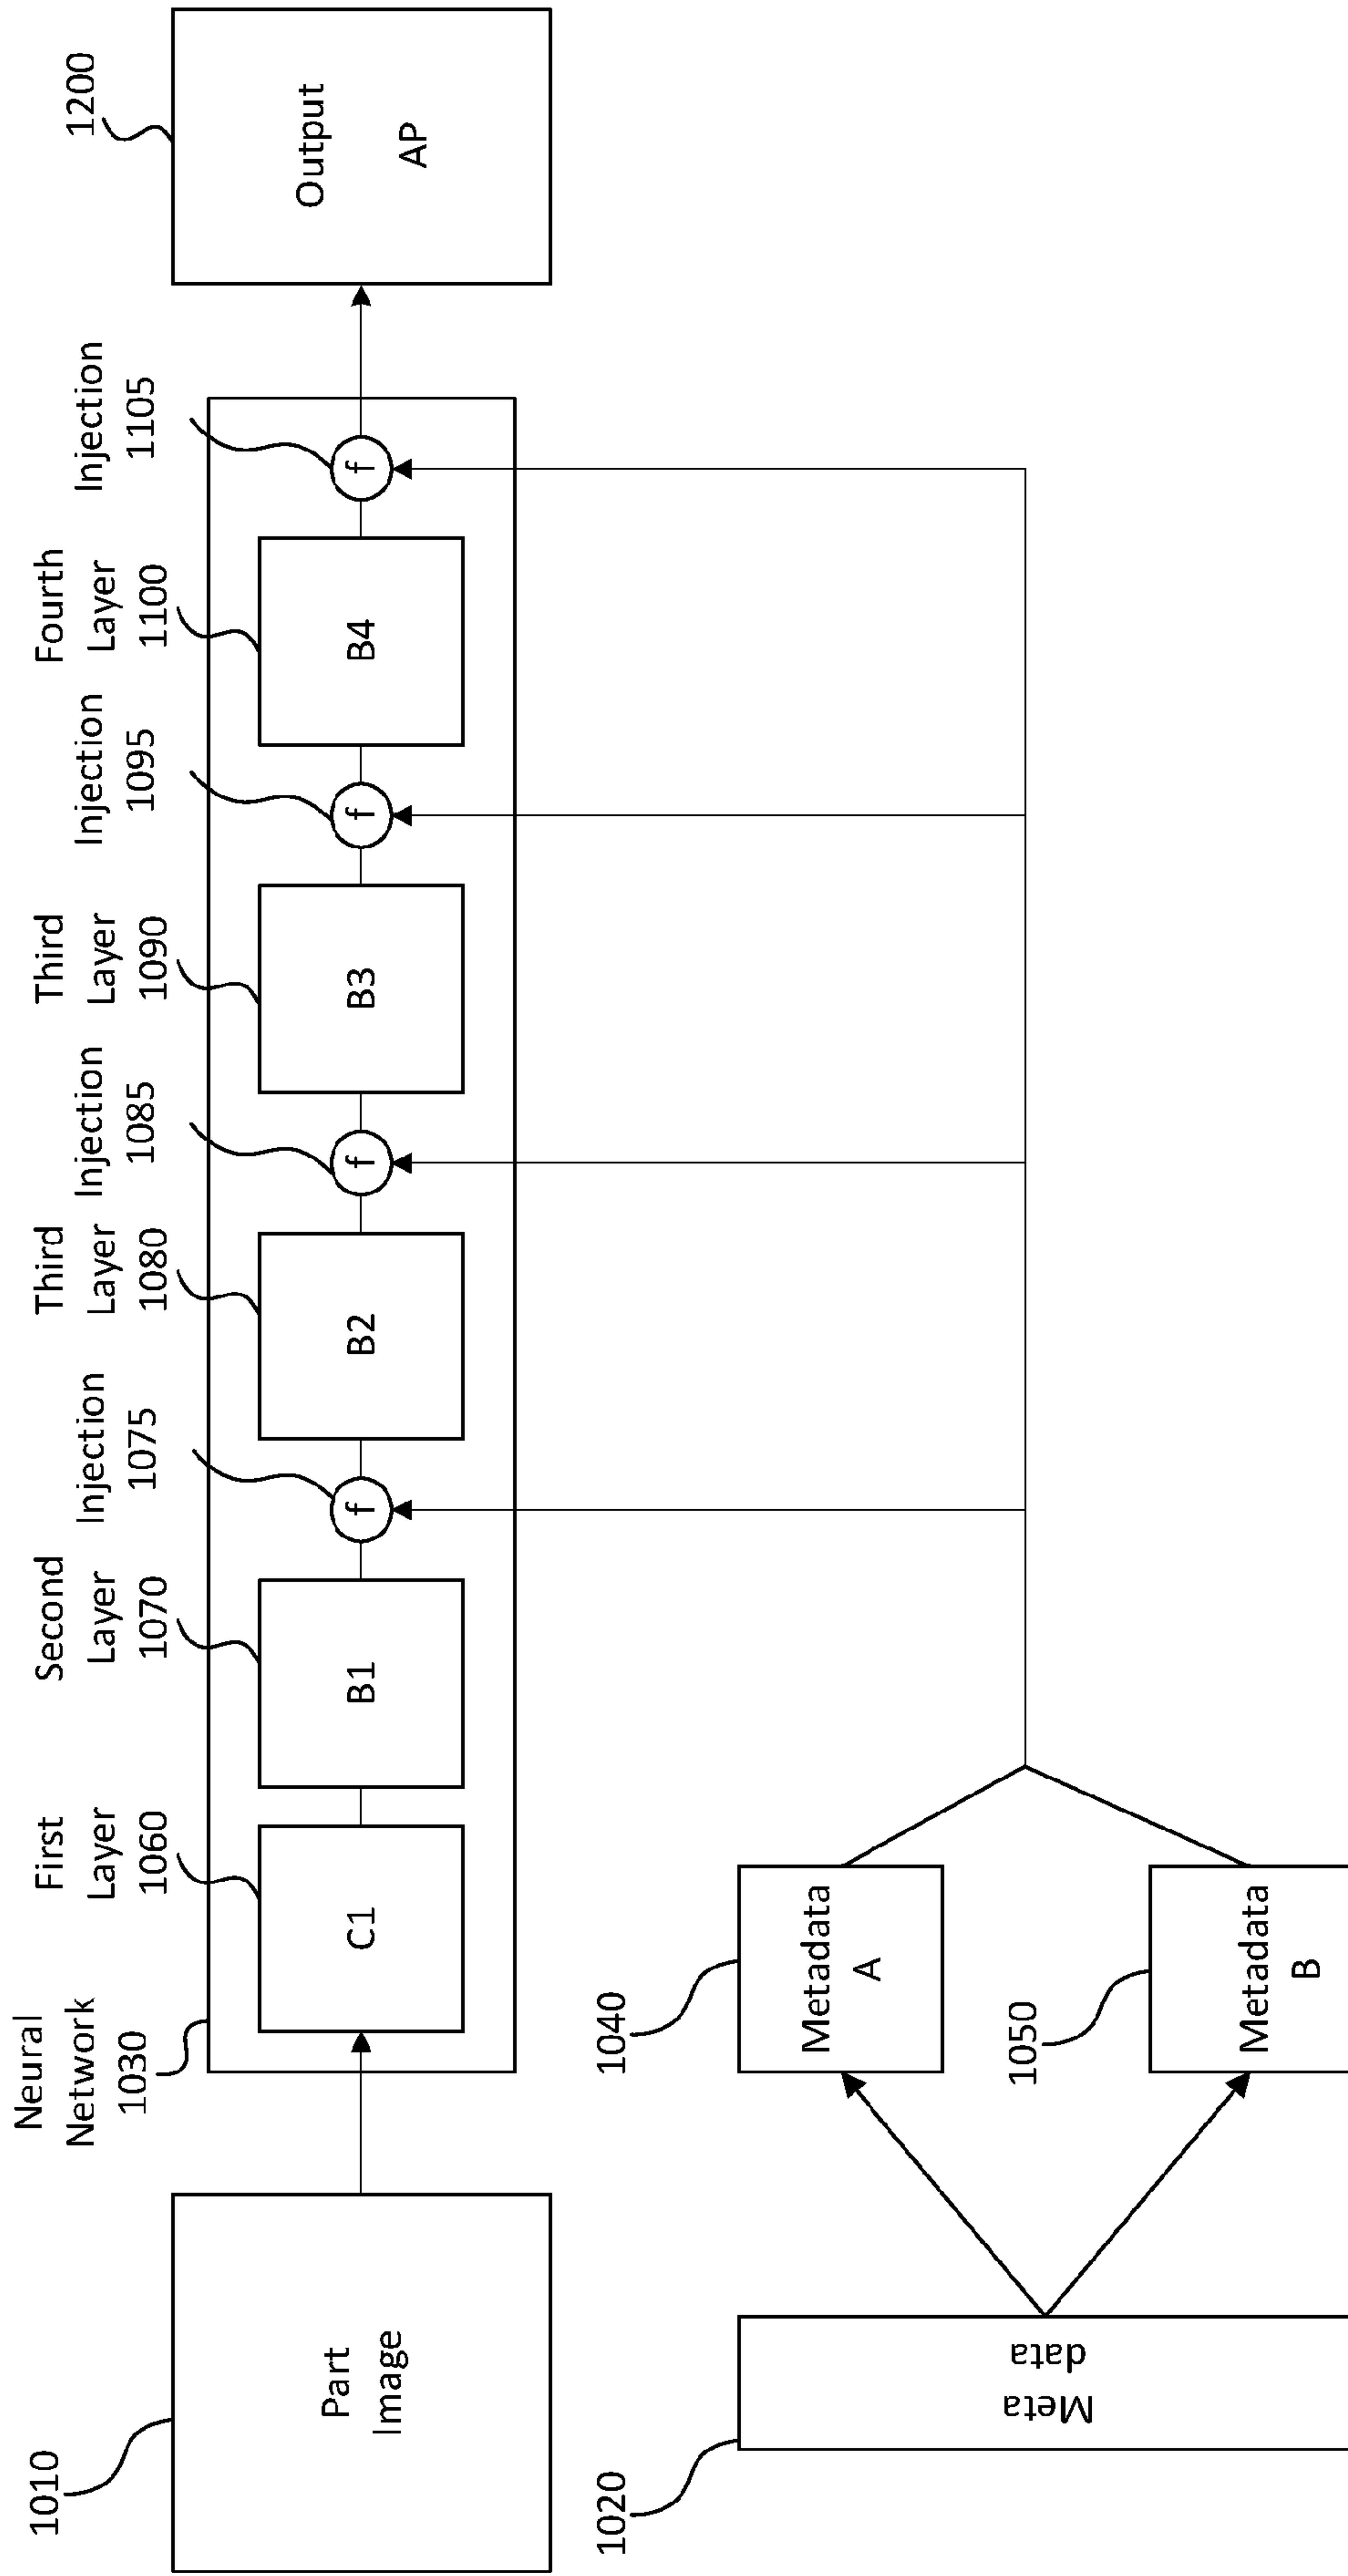


Figure 8



Estimate Repairs From Photos
900

Figure 9



Neural Network
Structure
1000

Figure 10




Figure 11

1200
Automated Vehicle Repair Assessment




Figure 12

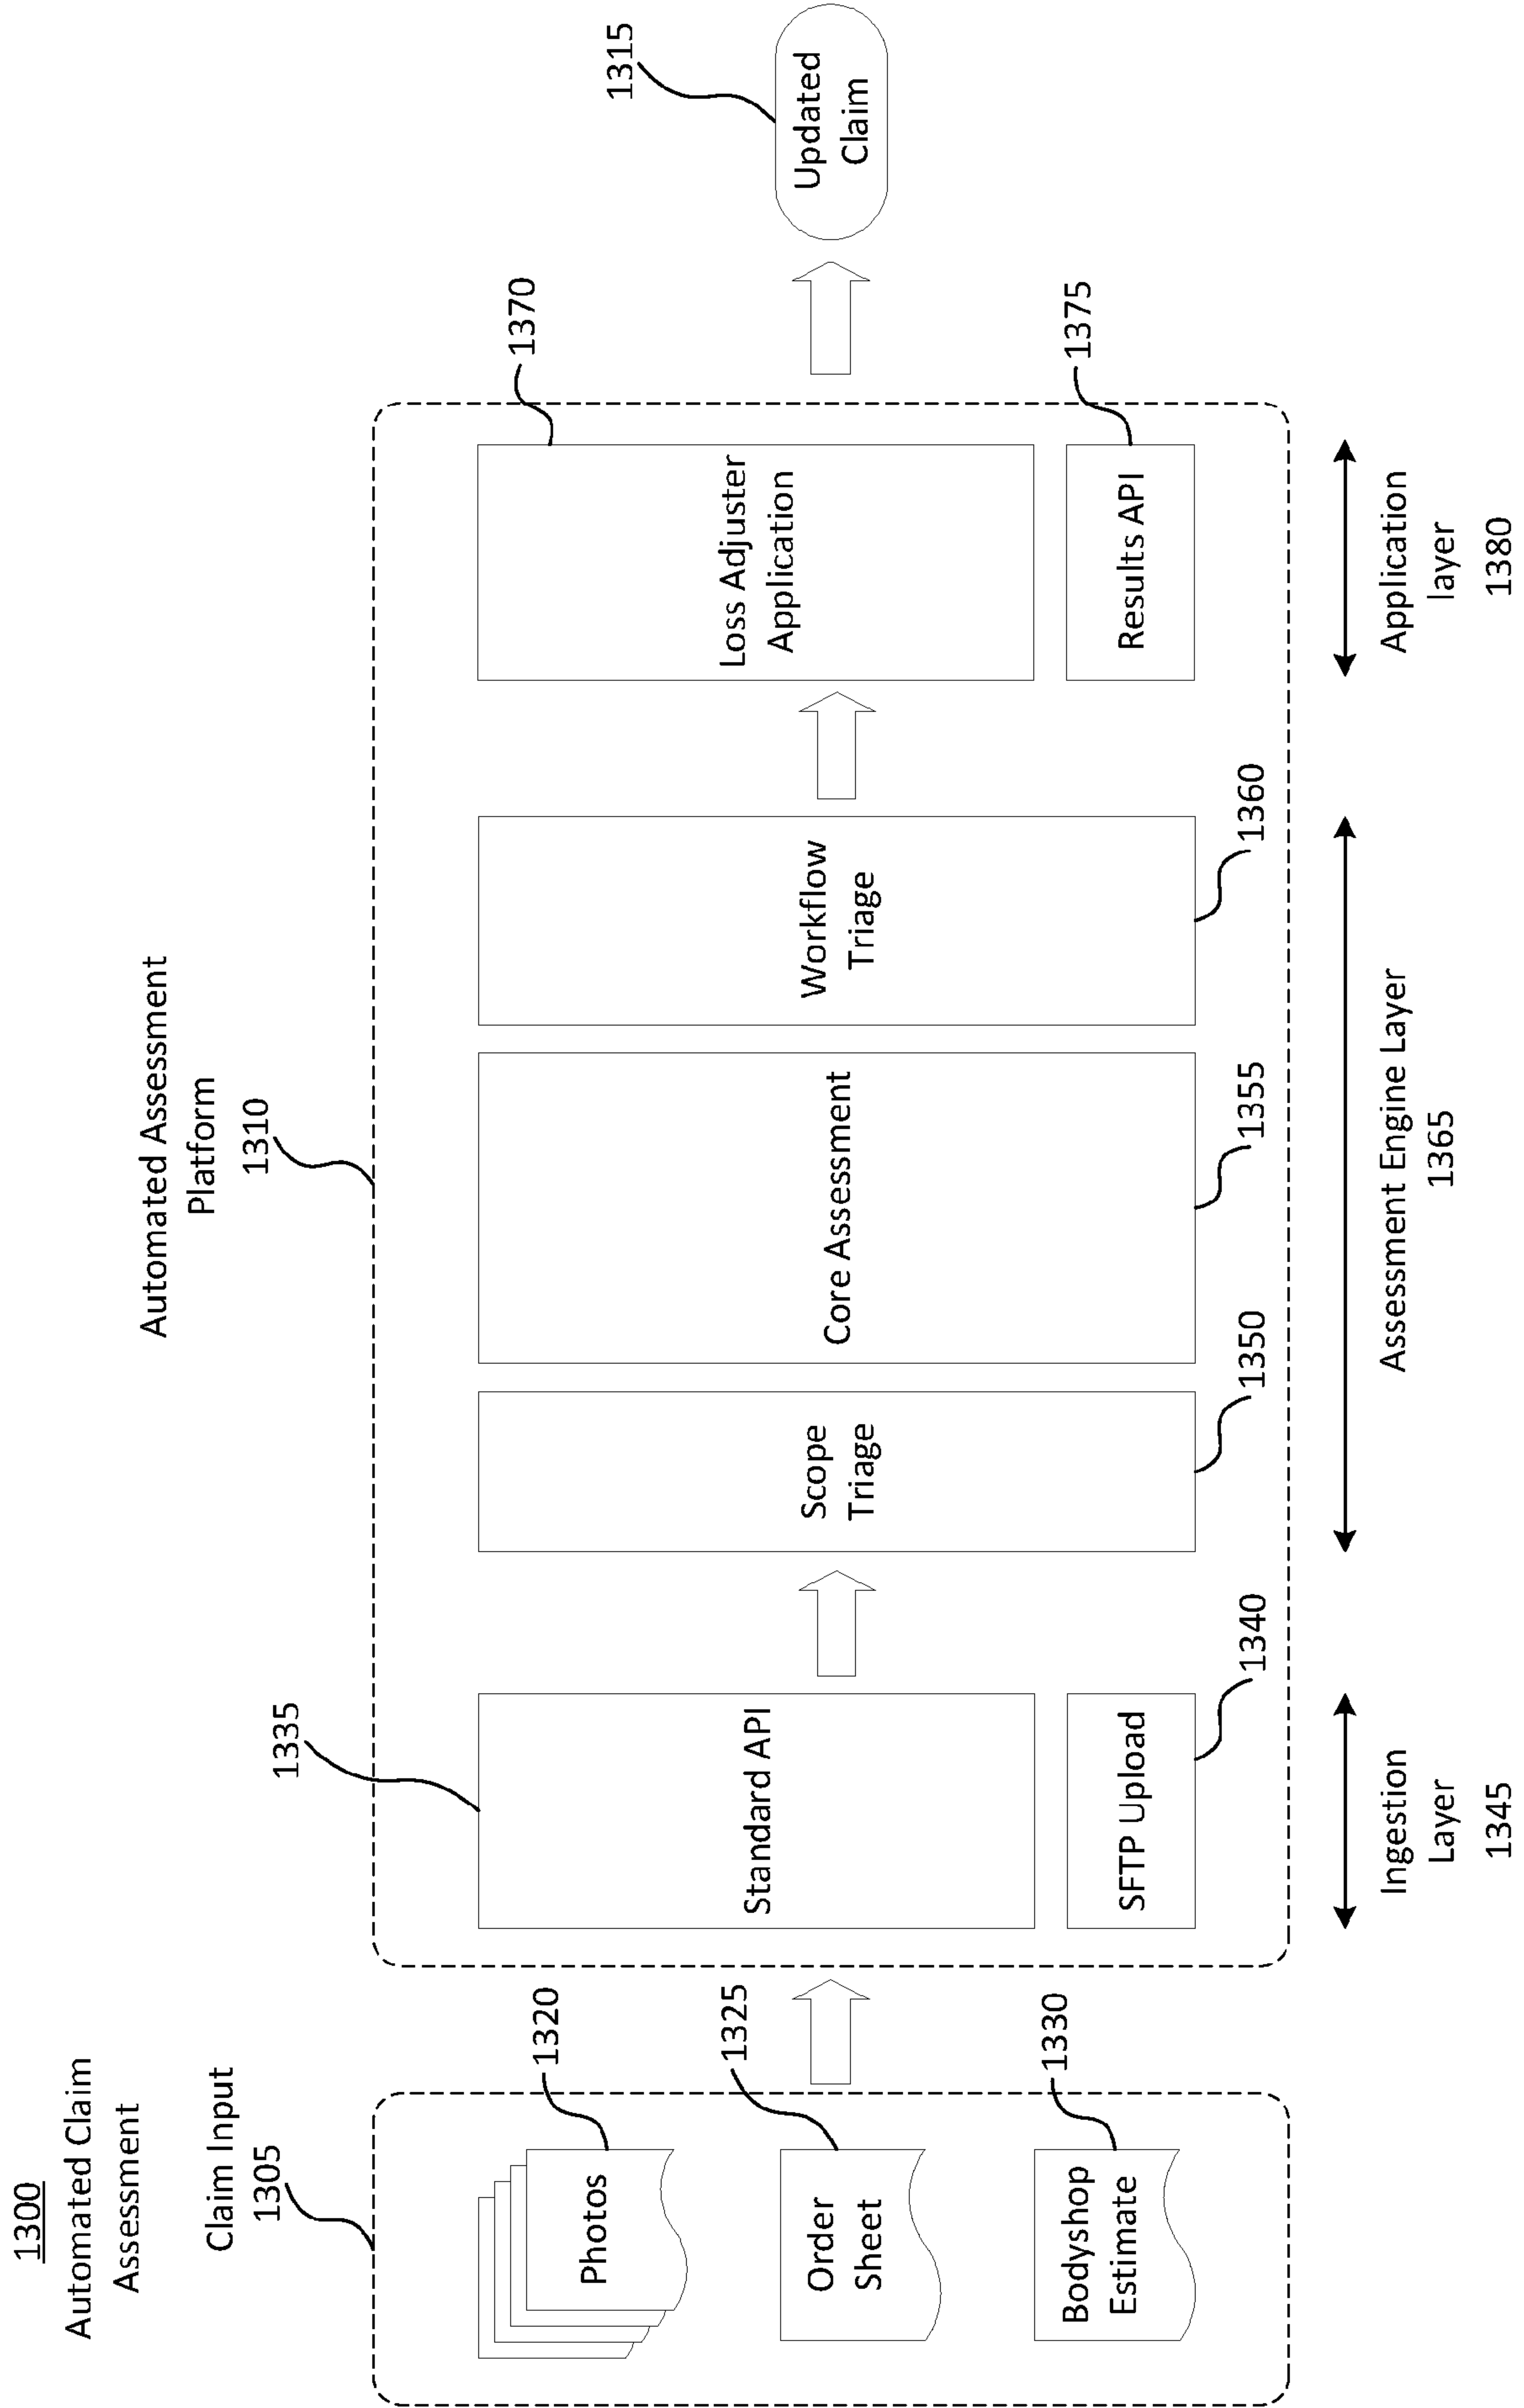


Figure 13




Figure 14

1500
Comparison Analysis Process

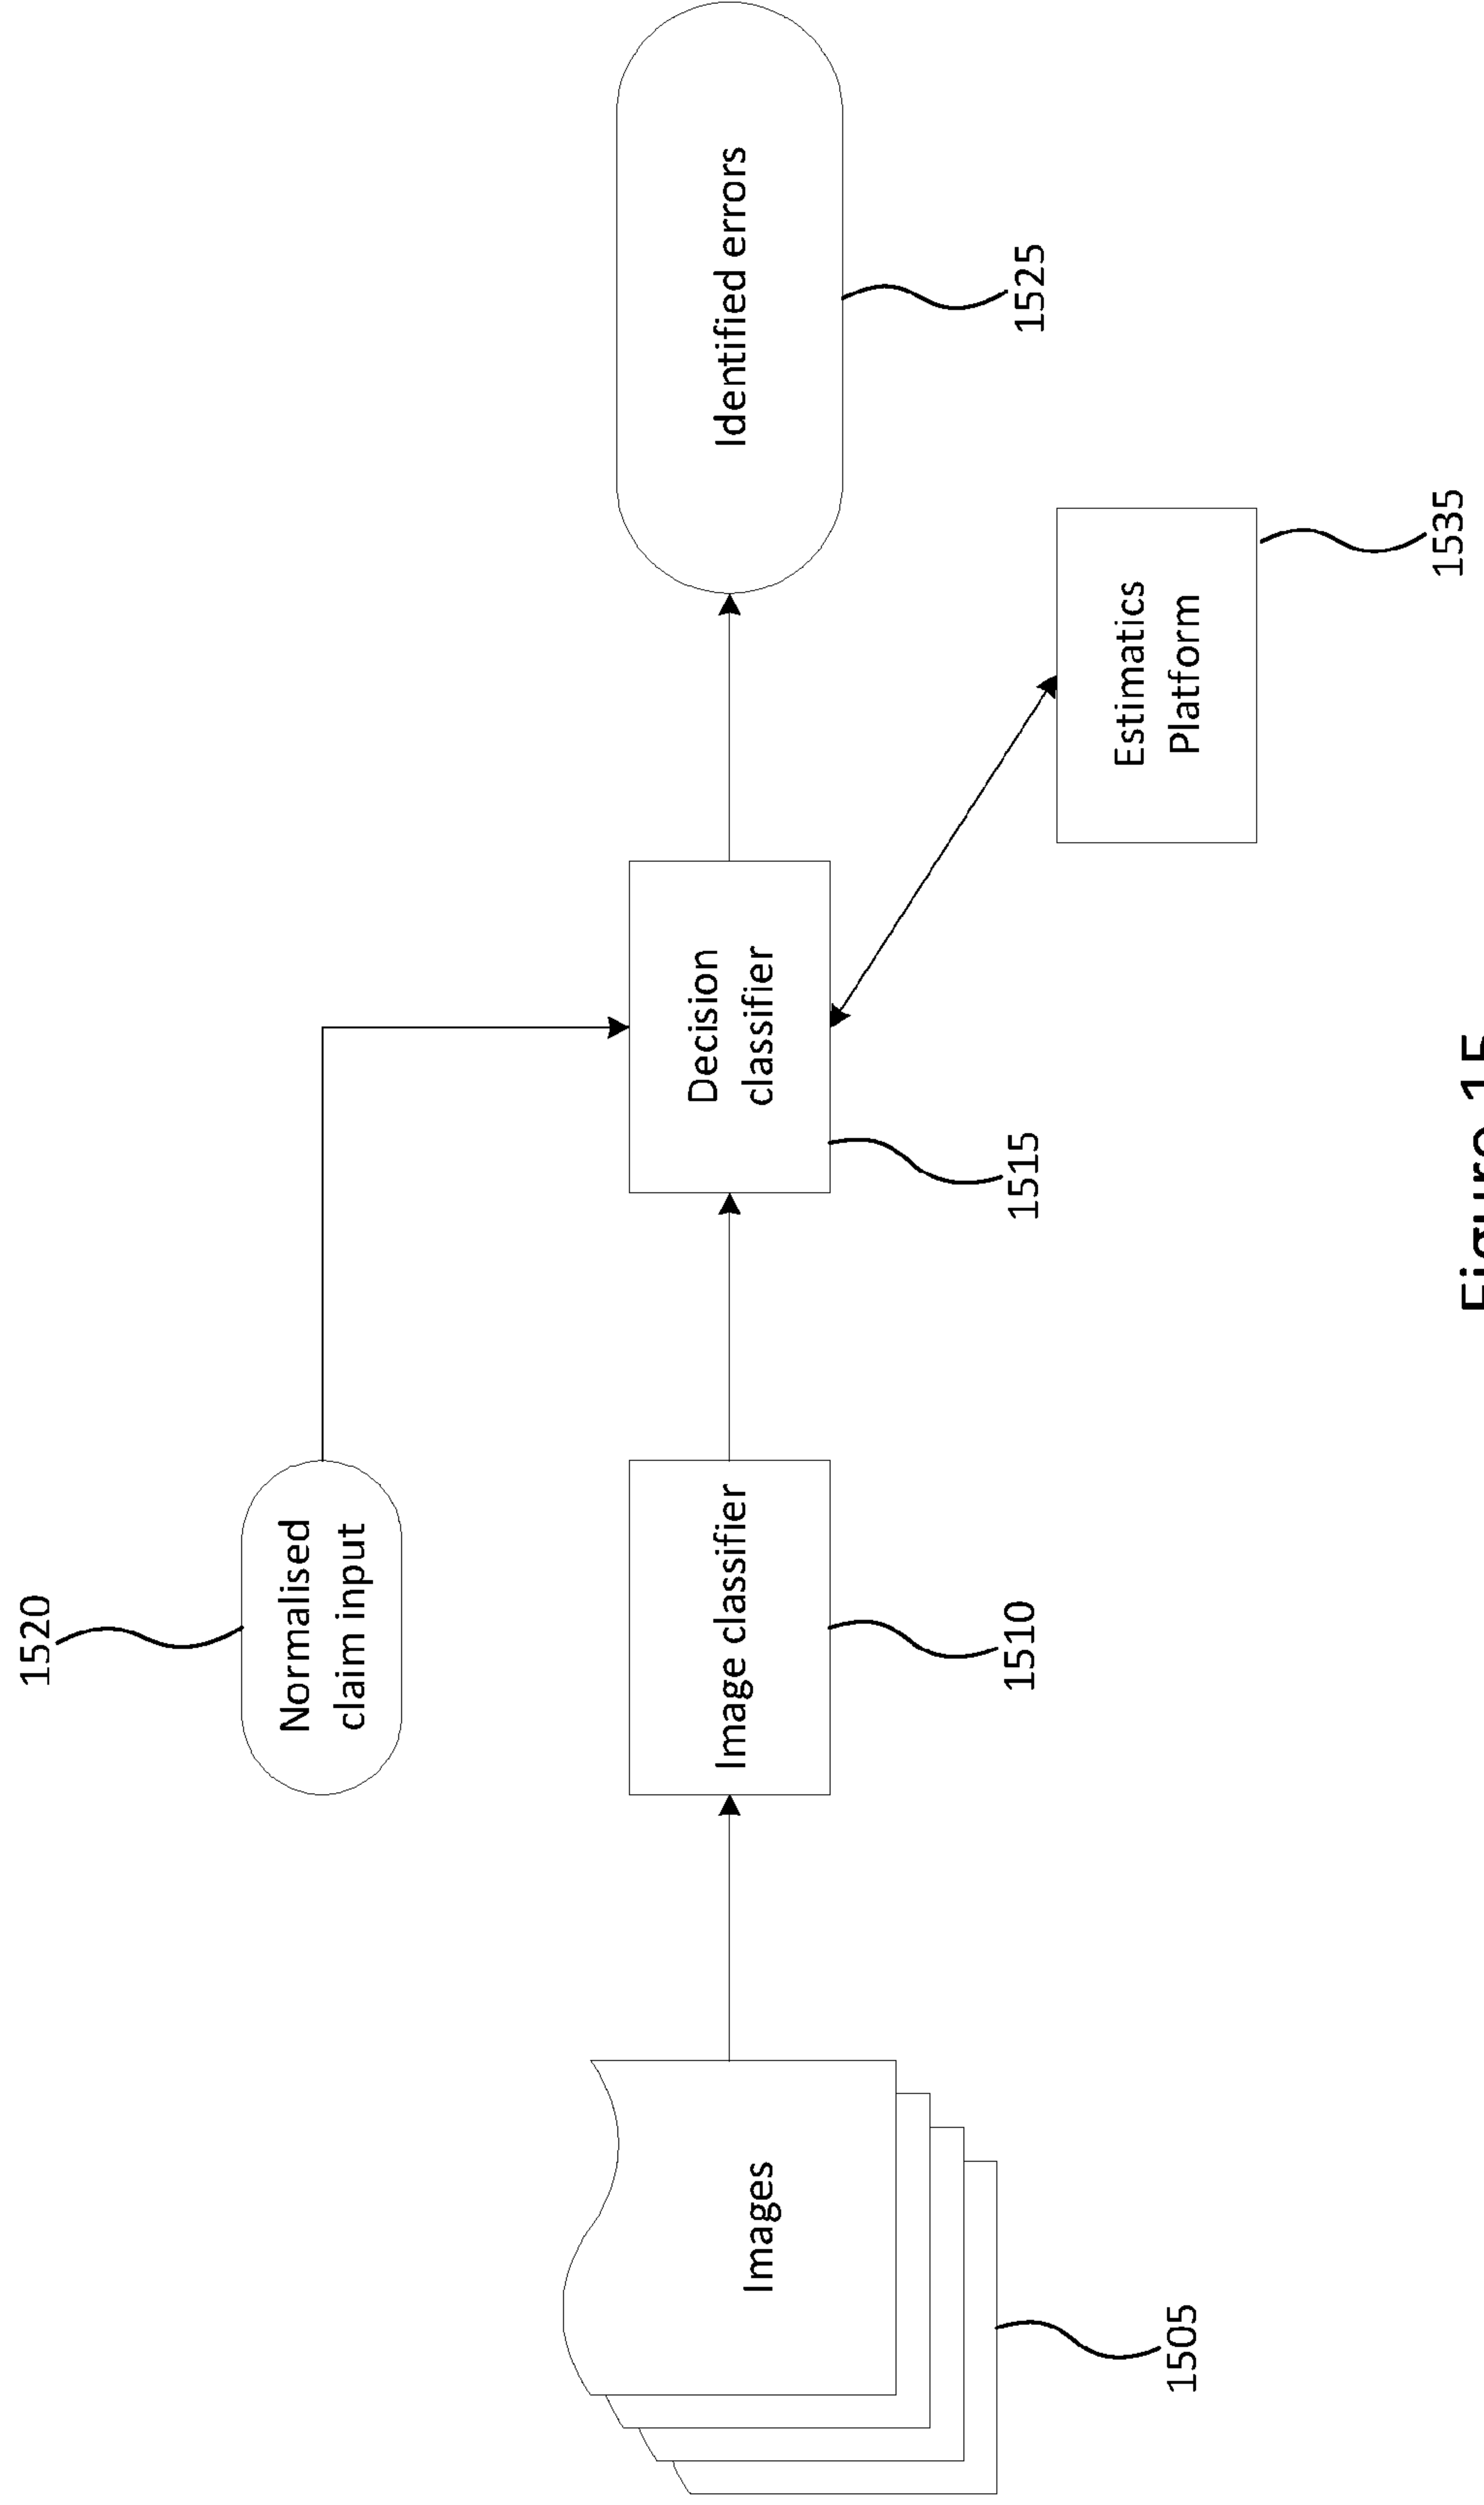


Figure 15

1600
Comparison Analysis Process

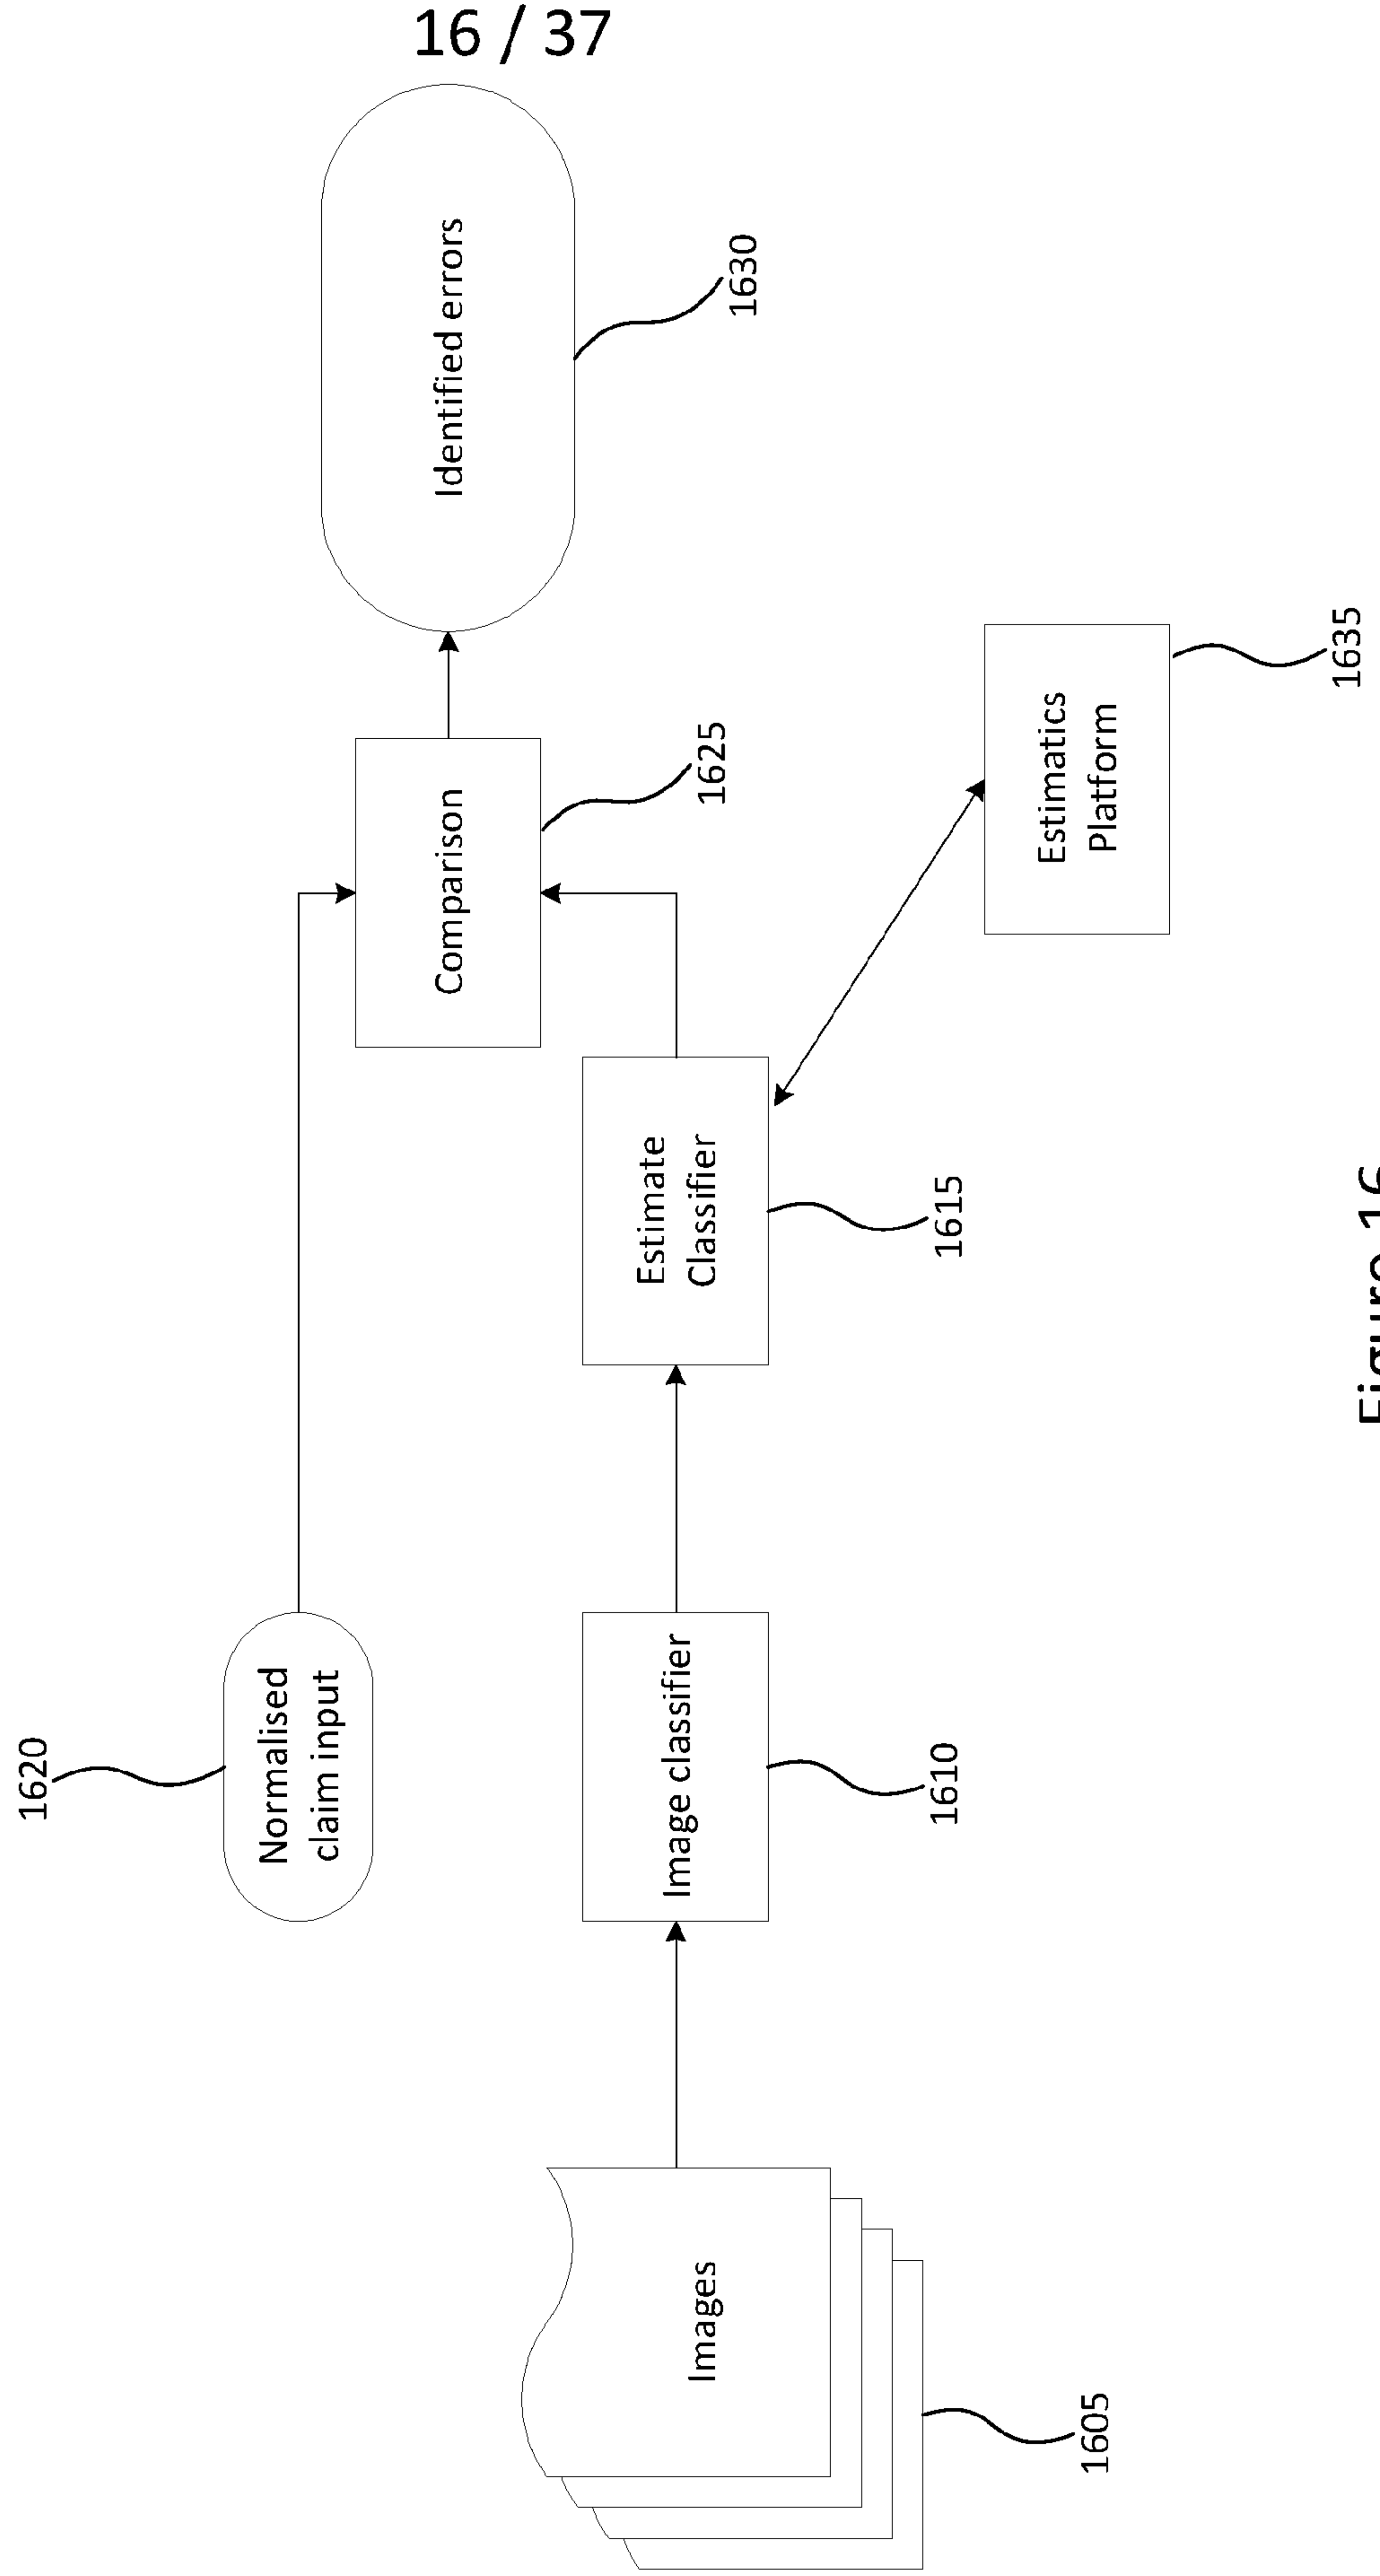


Figure 16

1700
Pre-Checked Vehicle Repair Claim

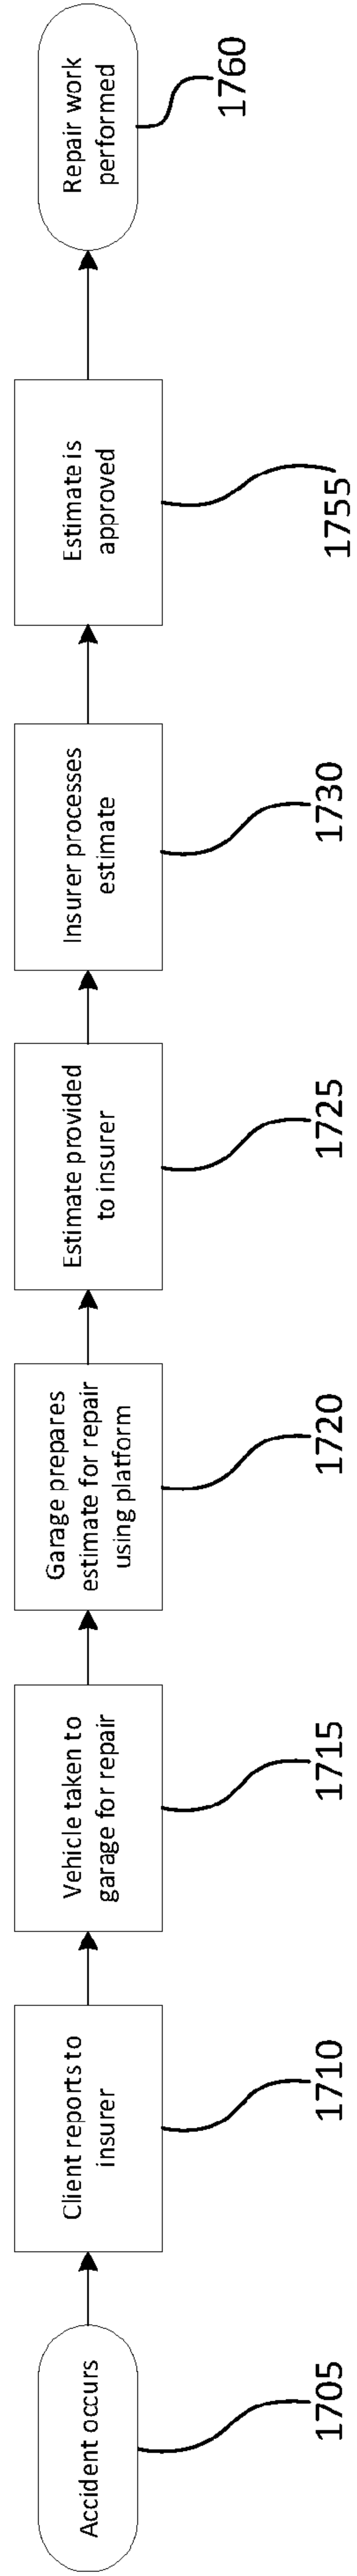


Figure 17

Generalised Method Overview

1800

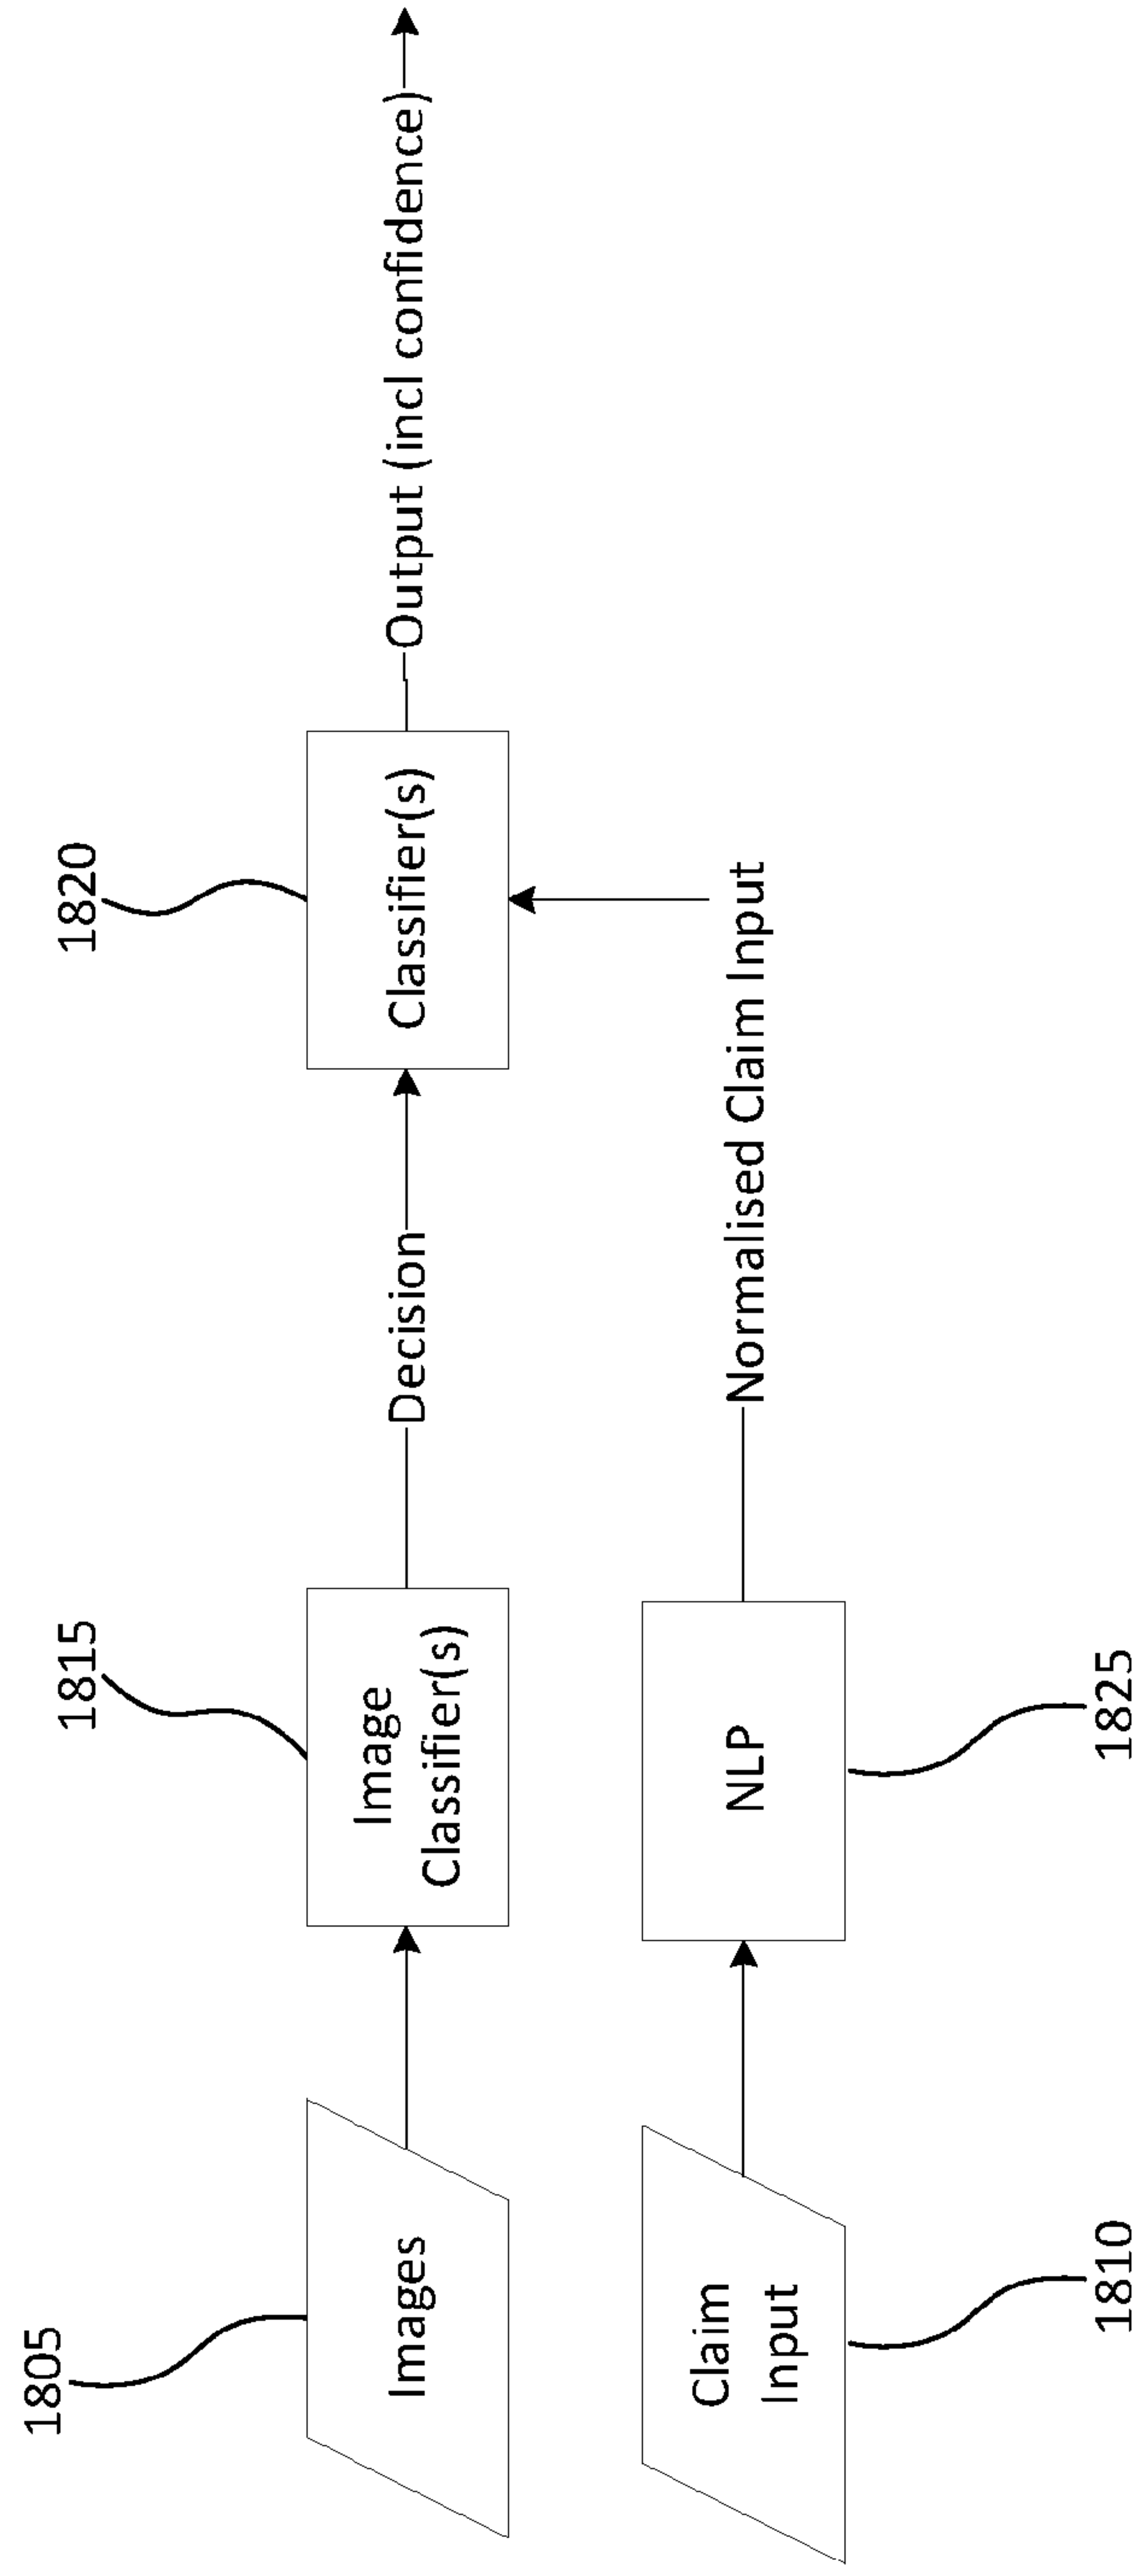


Figure 18

Normalised Car Parts

1900




Figure 19

Example Embodiment Detailed
Process Flowchart

2000

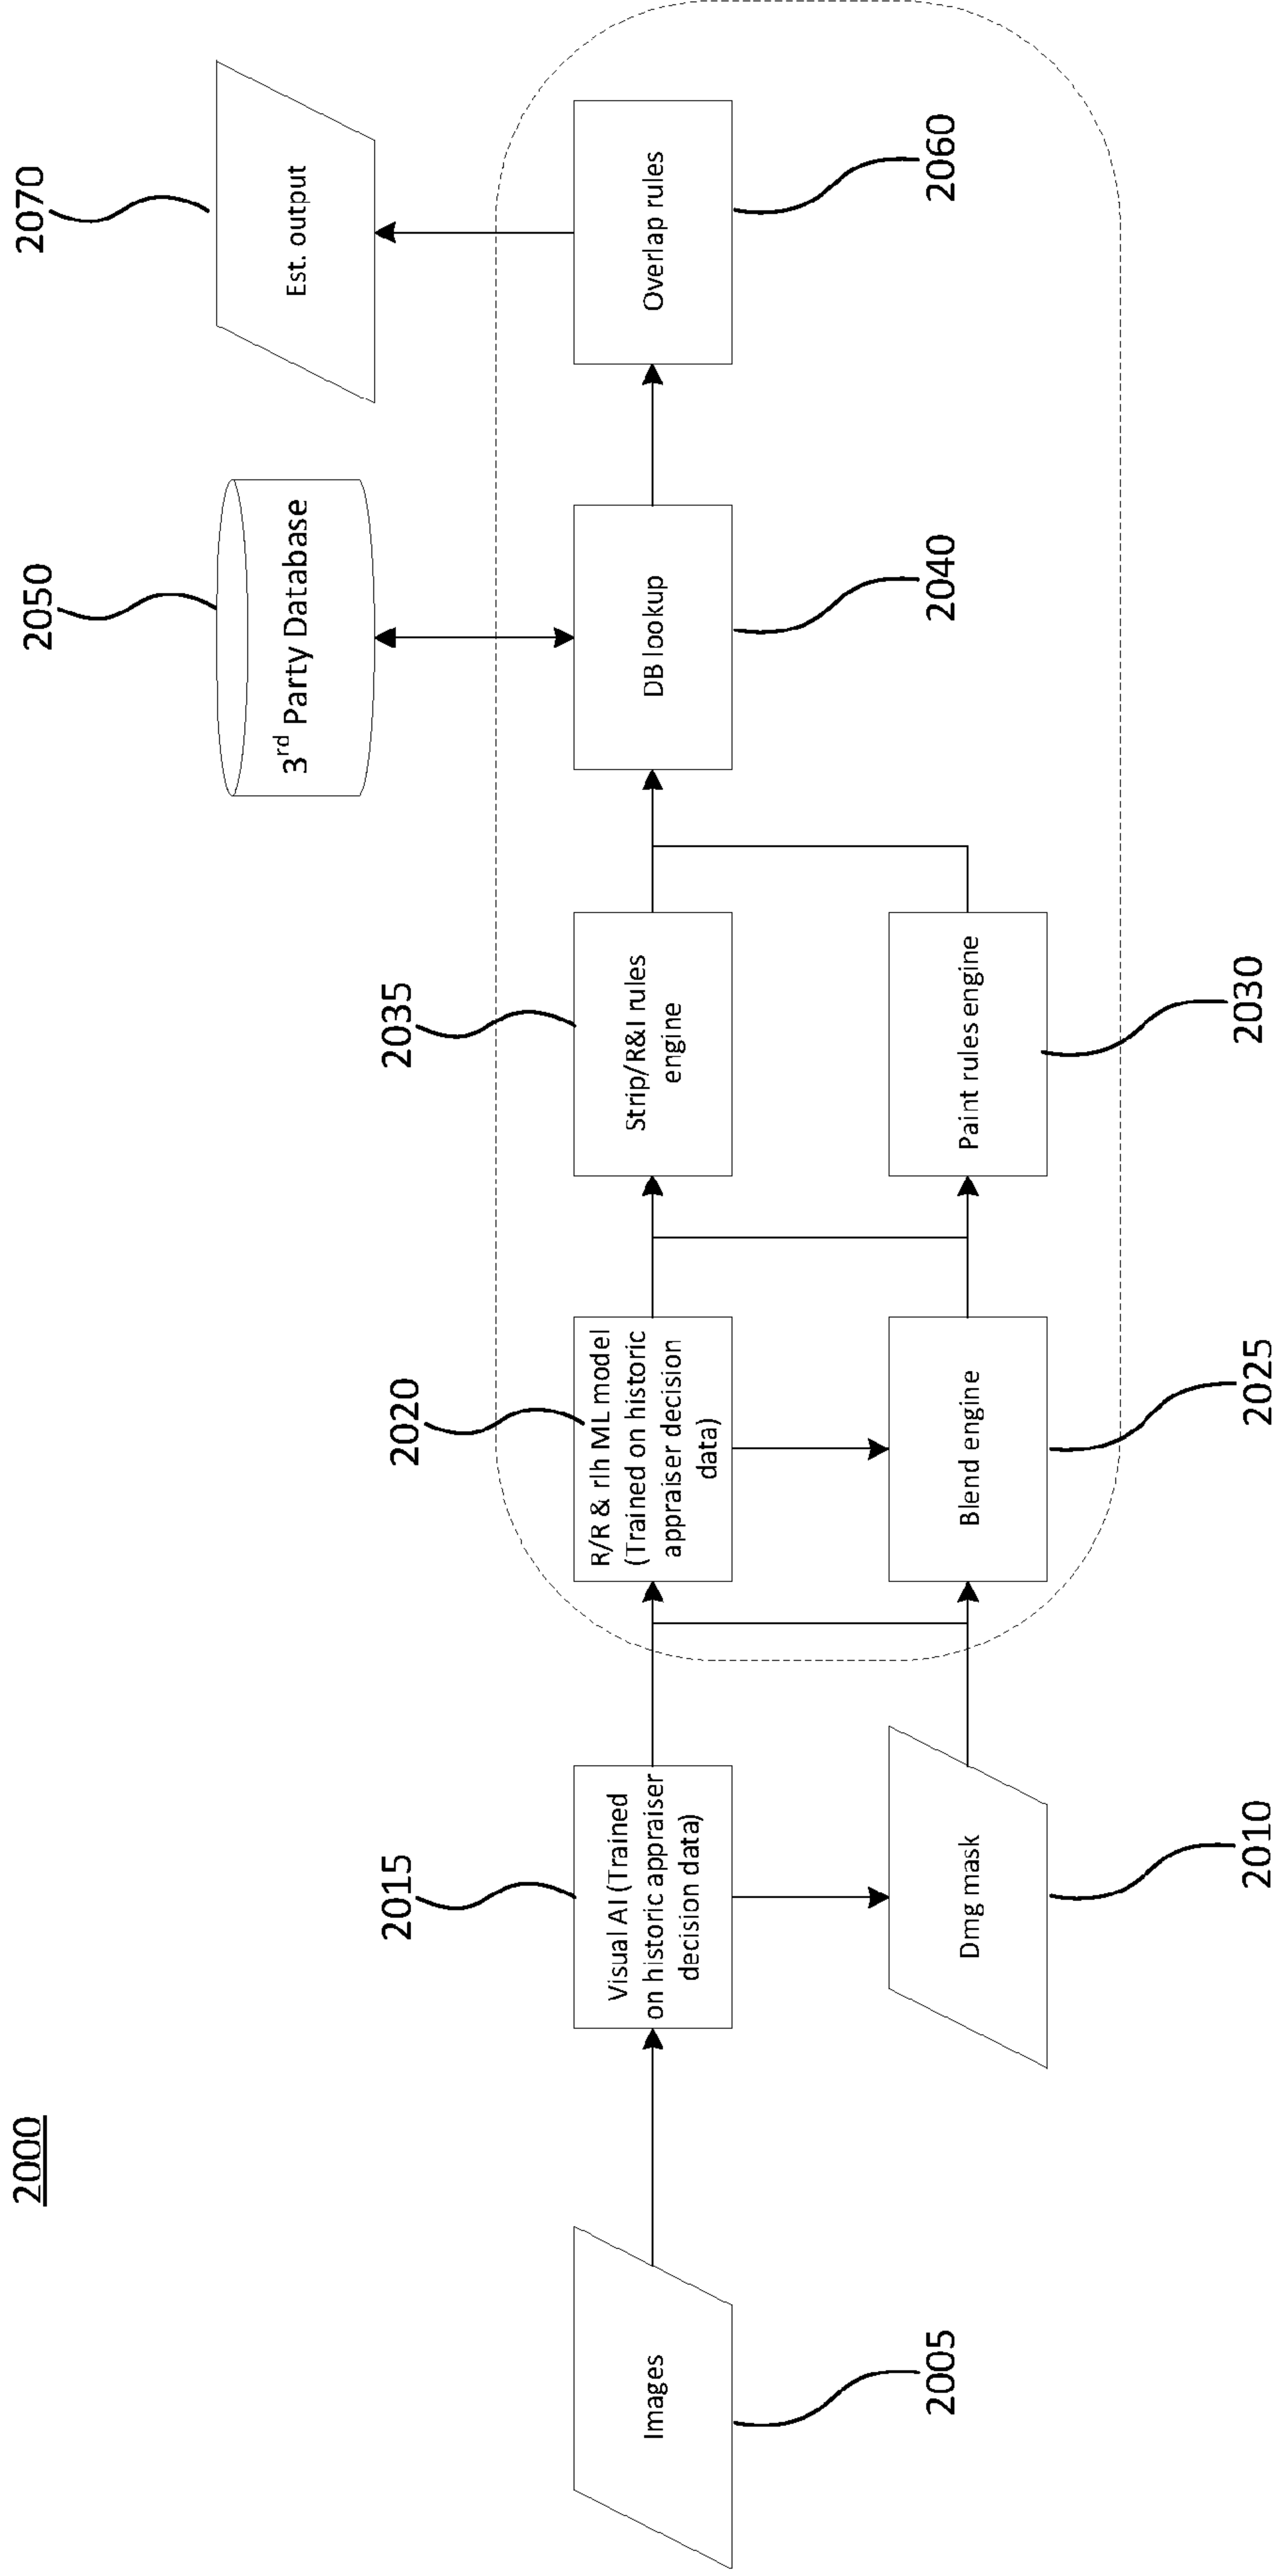


Figure 20

Normalised
Vehicle Parts
2100




Figure 21

Paint Check Process 2200

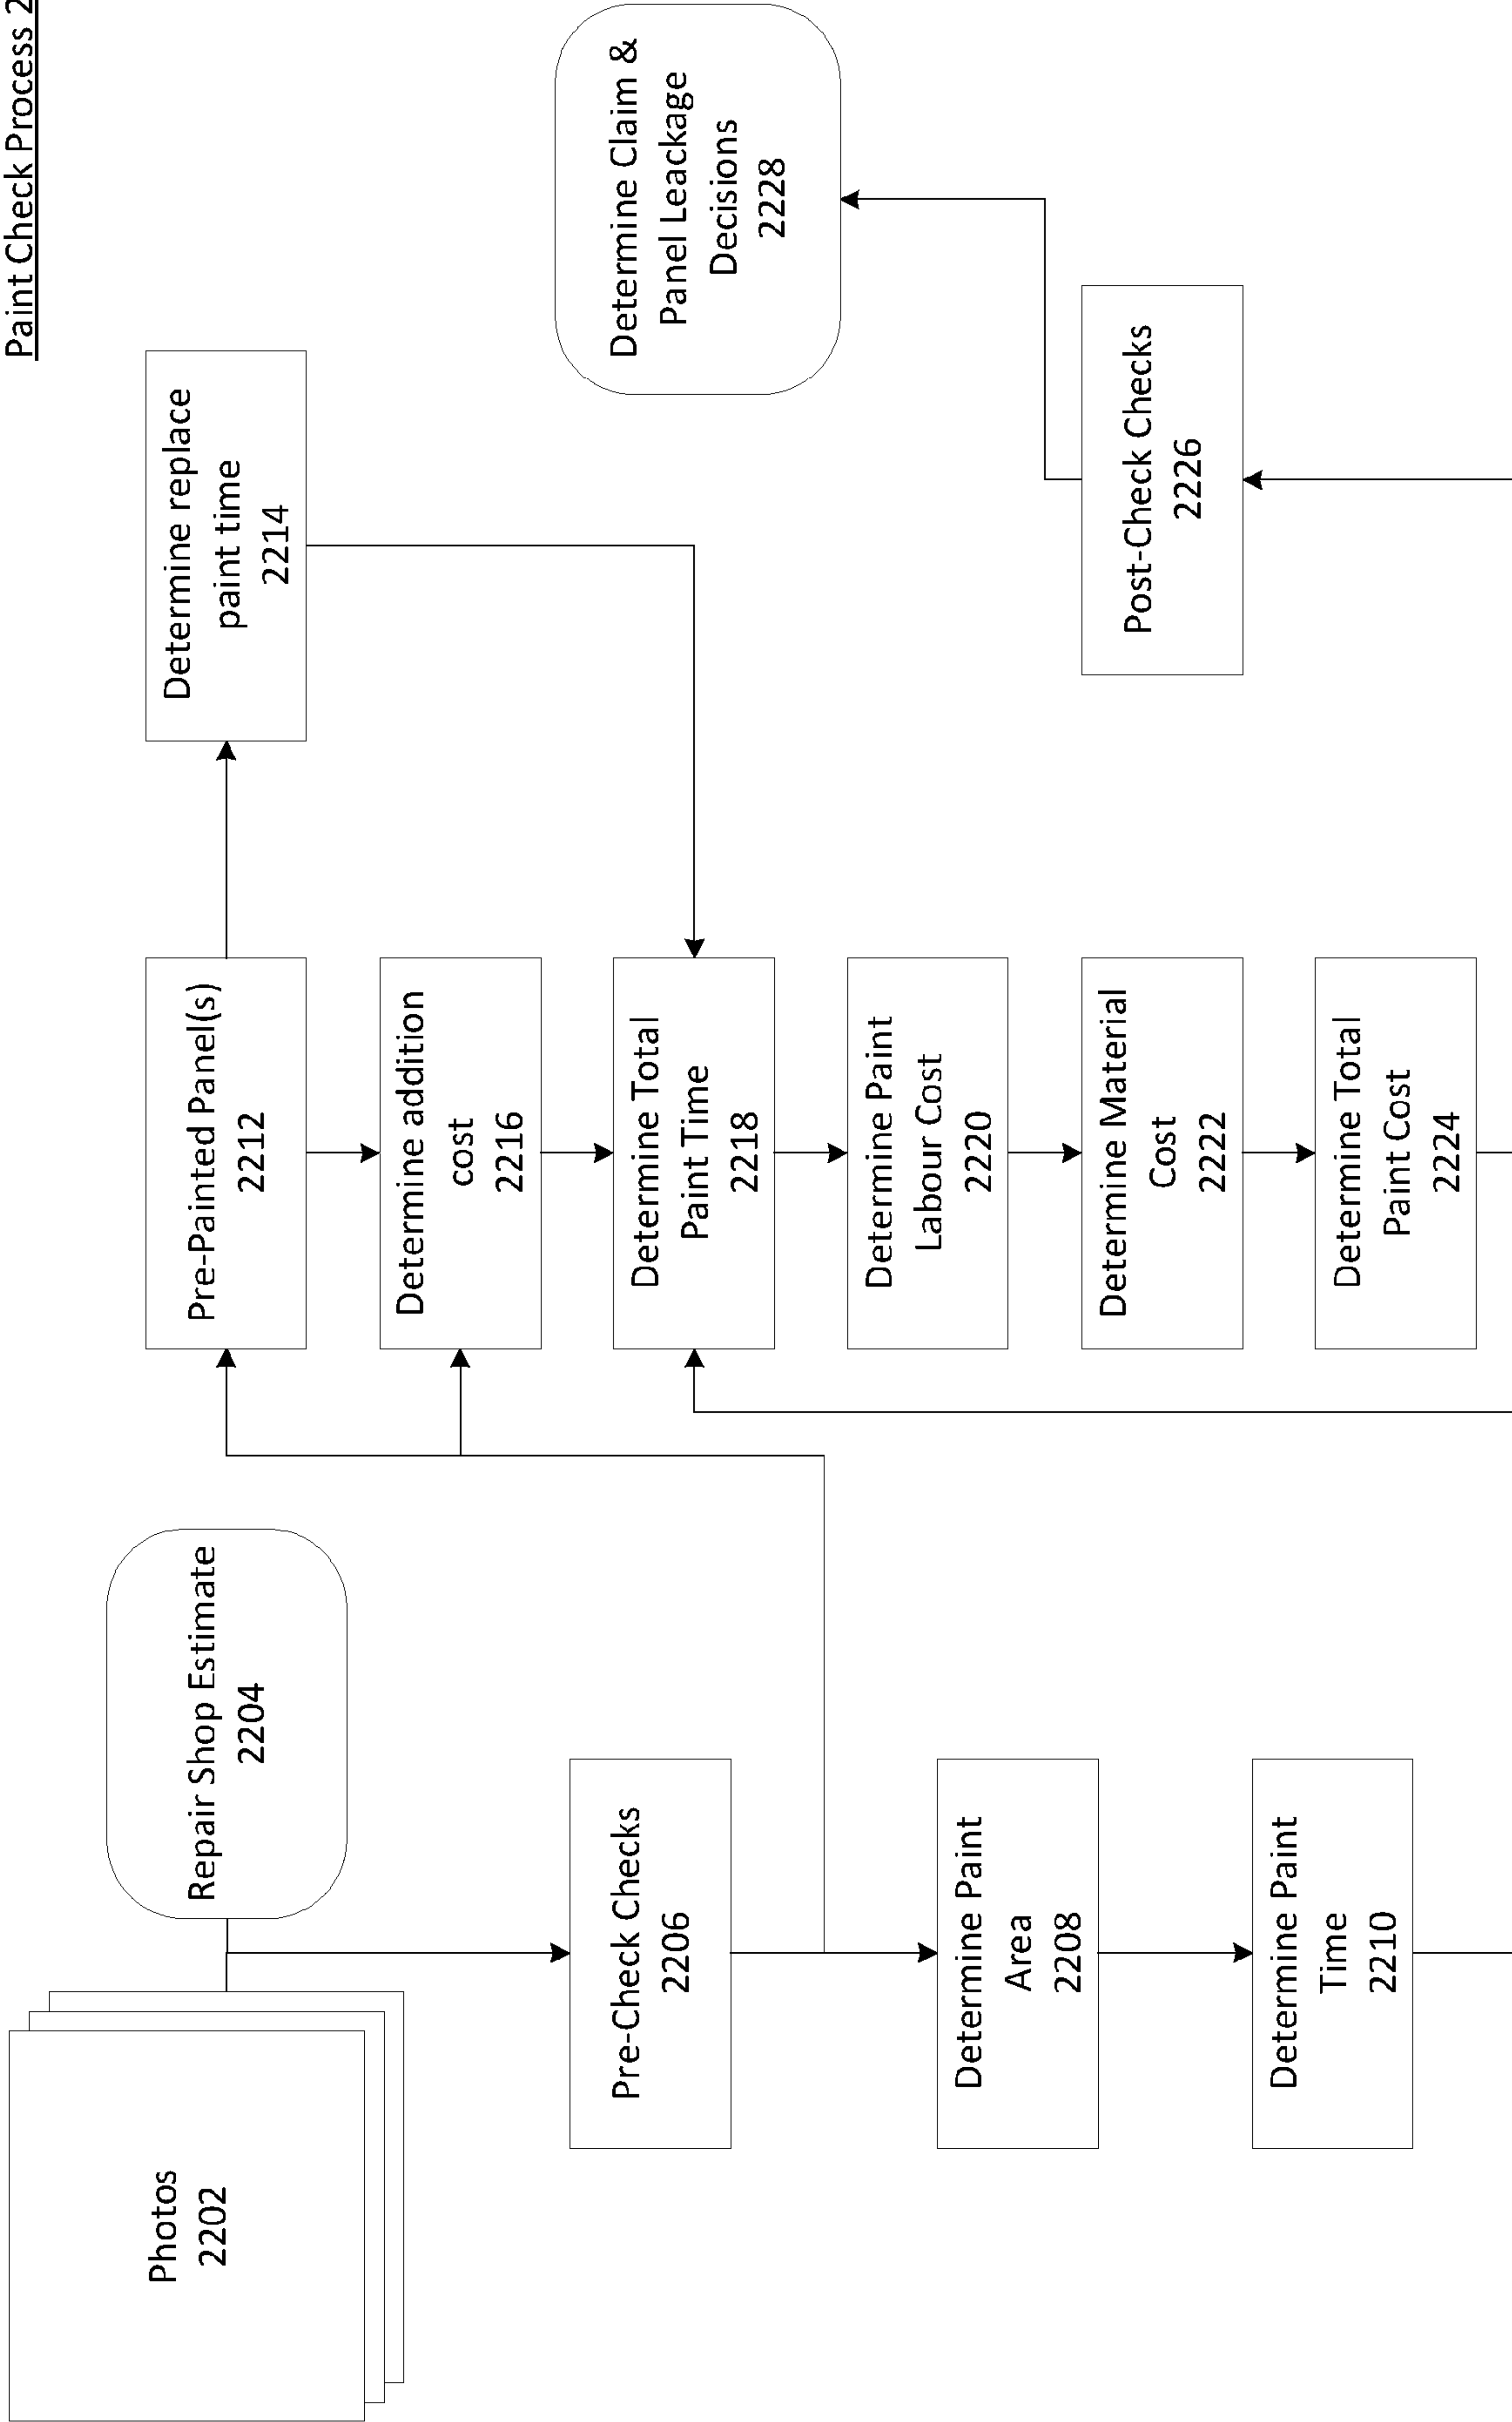


Figure 22

Jurisdiction Specific Rules 2300

Jurisdiction Specific Rules
2320

Decision on Paint Coverage Required
2310

Rules Specific to Country A
2322

Rules Specific to Country B
2324

Rules Specific to Country C
2326

Rules Specific to Country D
2328

Rules Specific to Country E
2330

	A	B	C	D	E
AI Decision					
Spot	1/5	Spot	Spot	Spot	30% Damaged
Minor	1/2	Less than 50%	20% or less	50% reduction in cost	60% Damaged
Major	1	More than 50%	More than 20%	20% reduction in cost	90% Damaged

Permitted Paint Coverage Amounts
2330

Figure 23

Model Architecture 2400

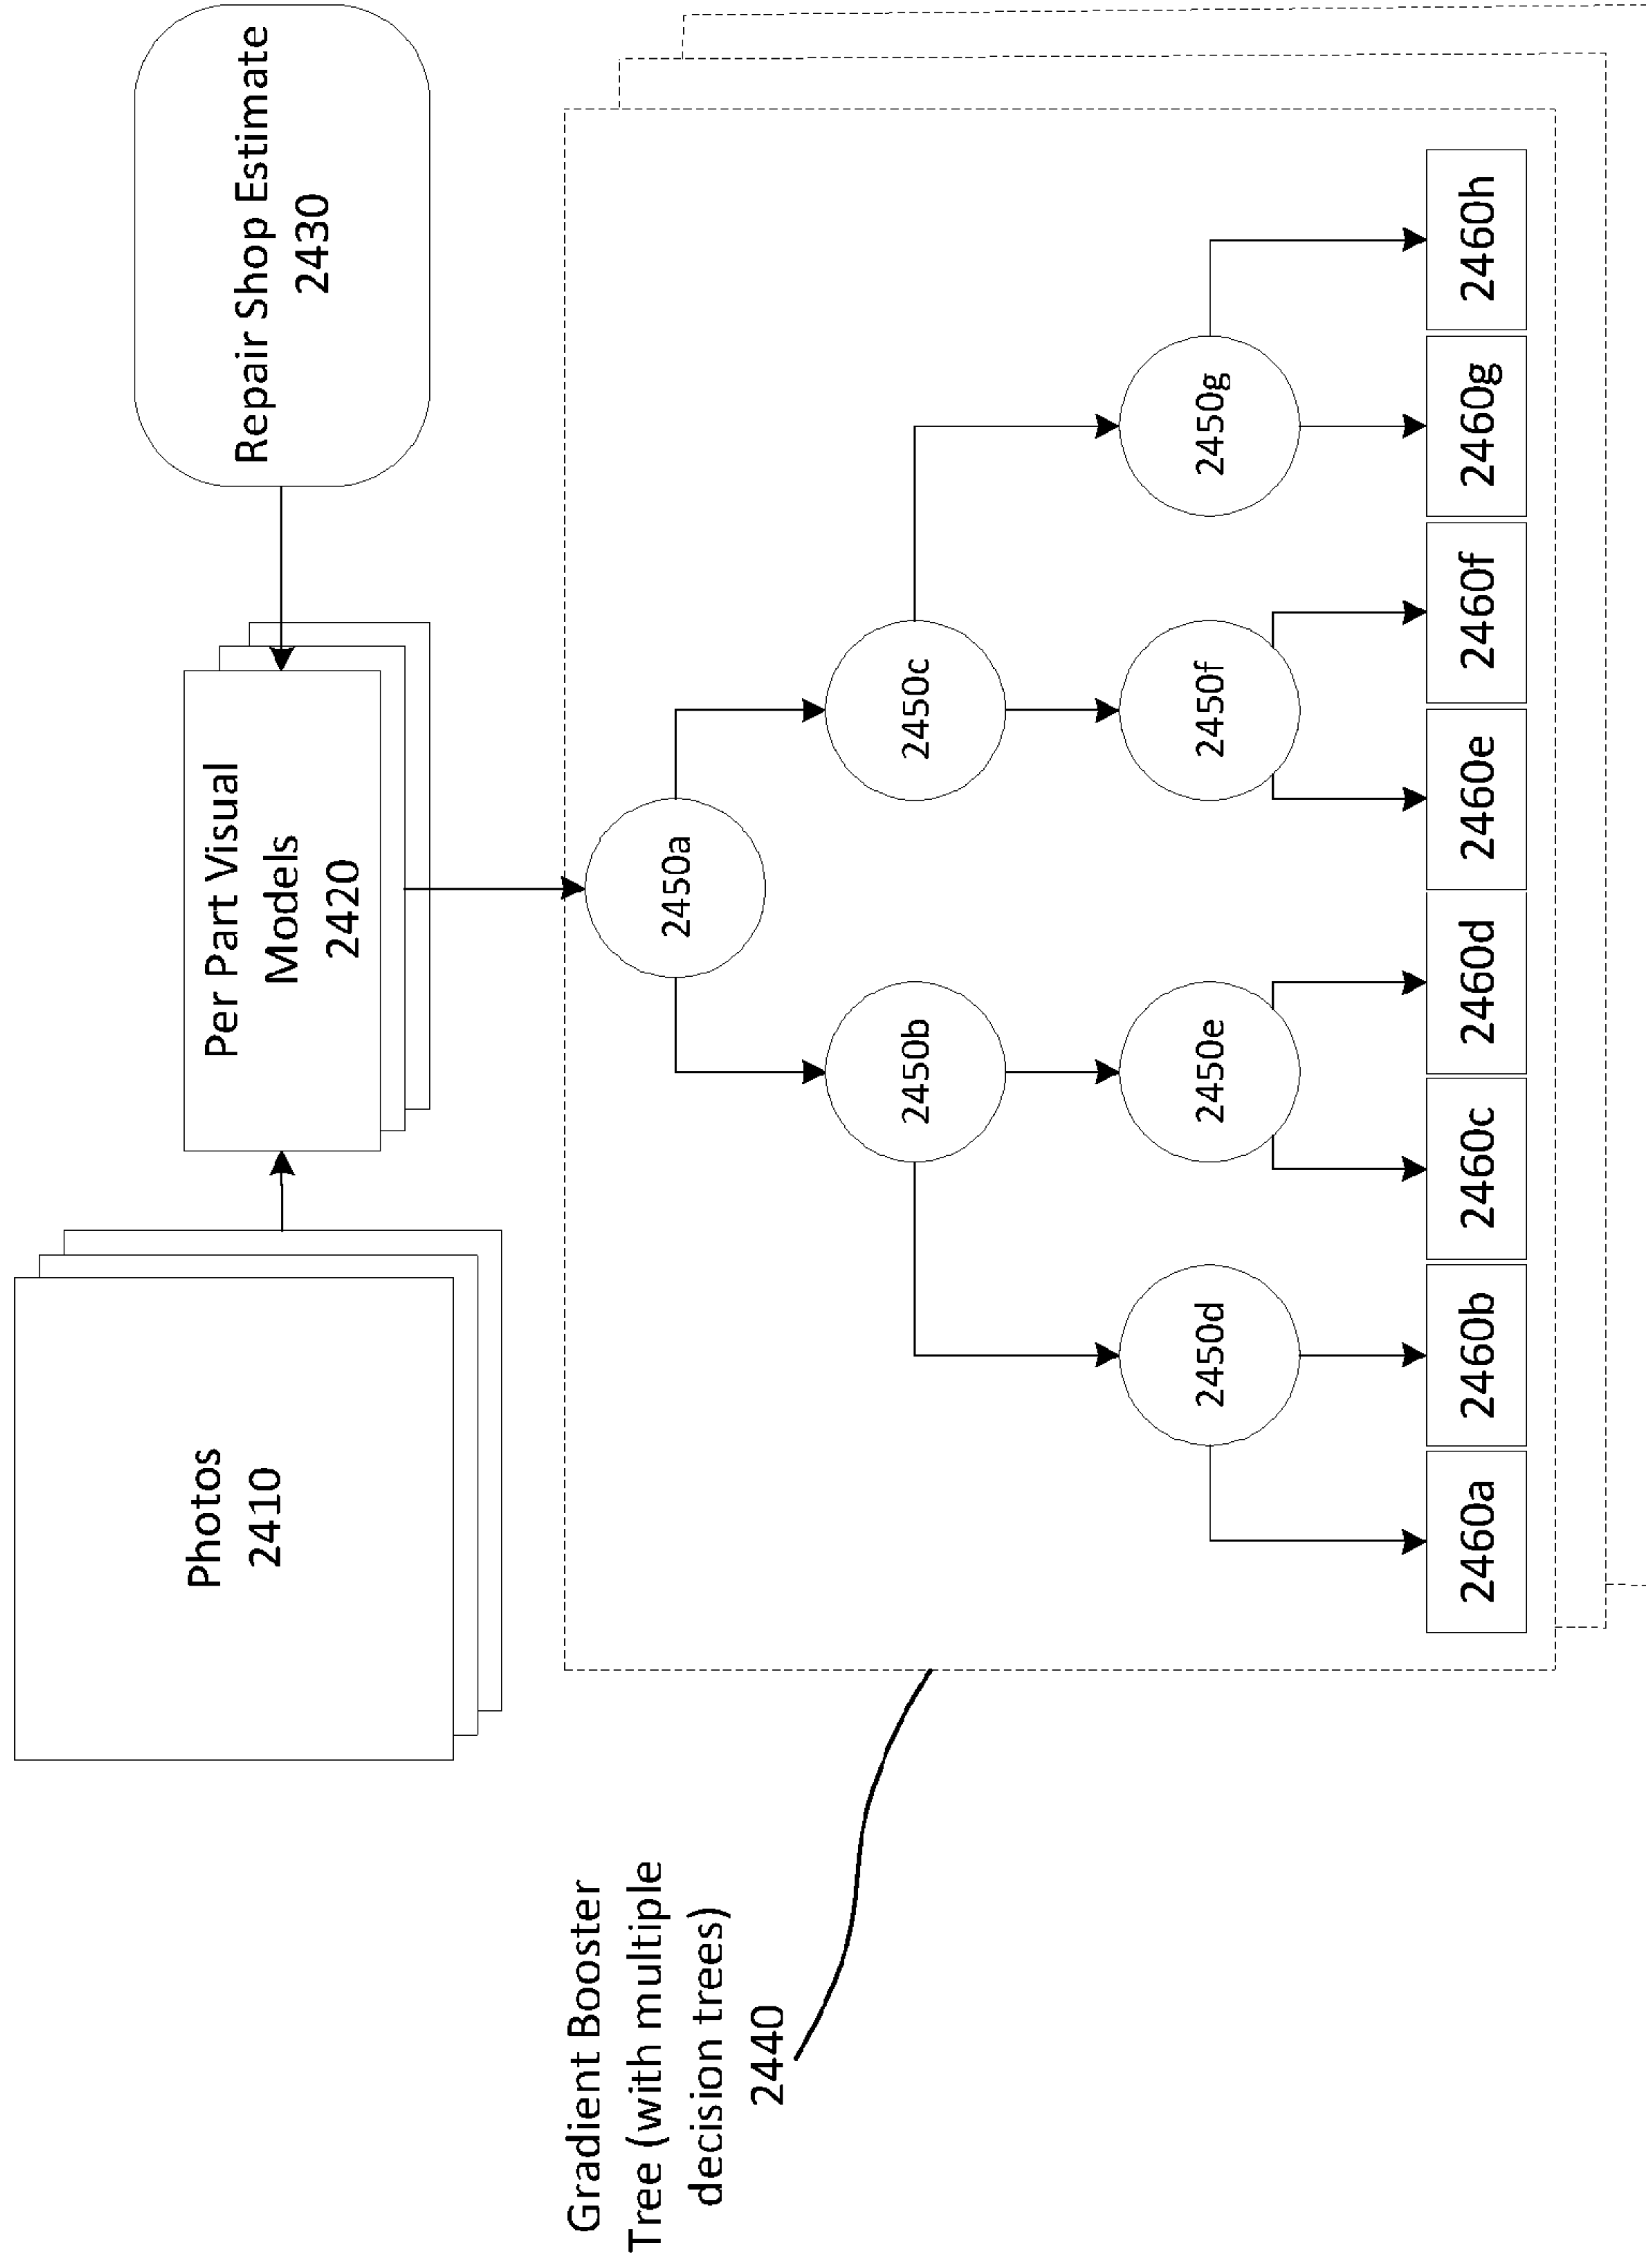


Figure 24

Calculating Paint Cost 2500

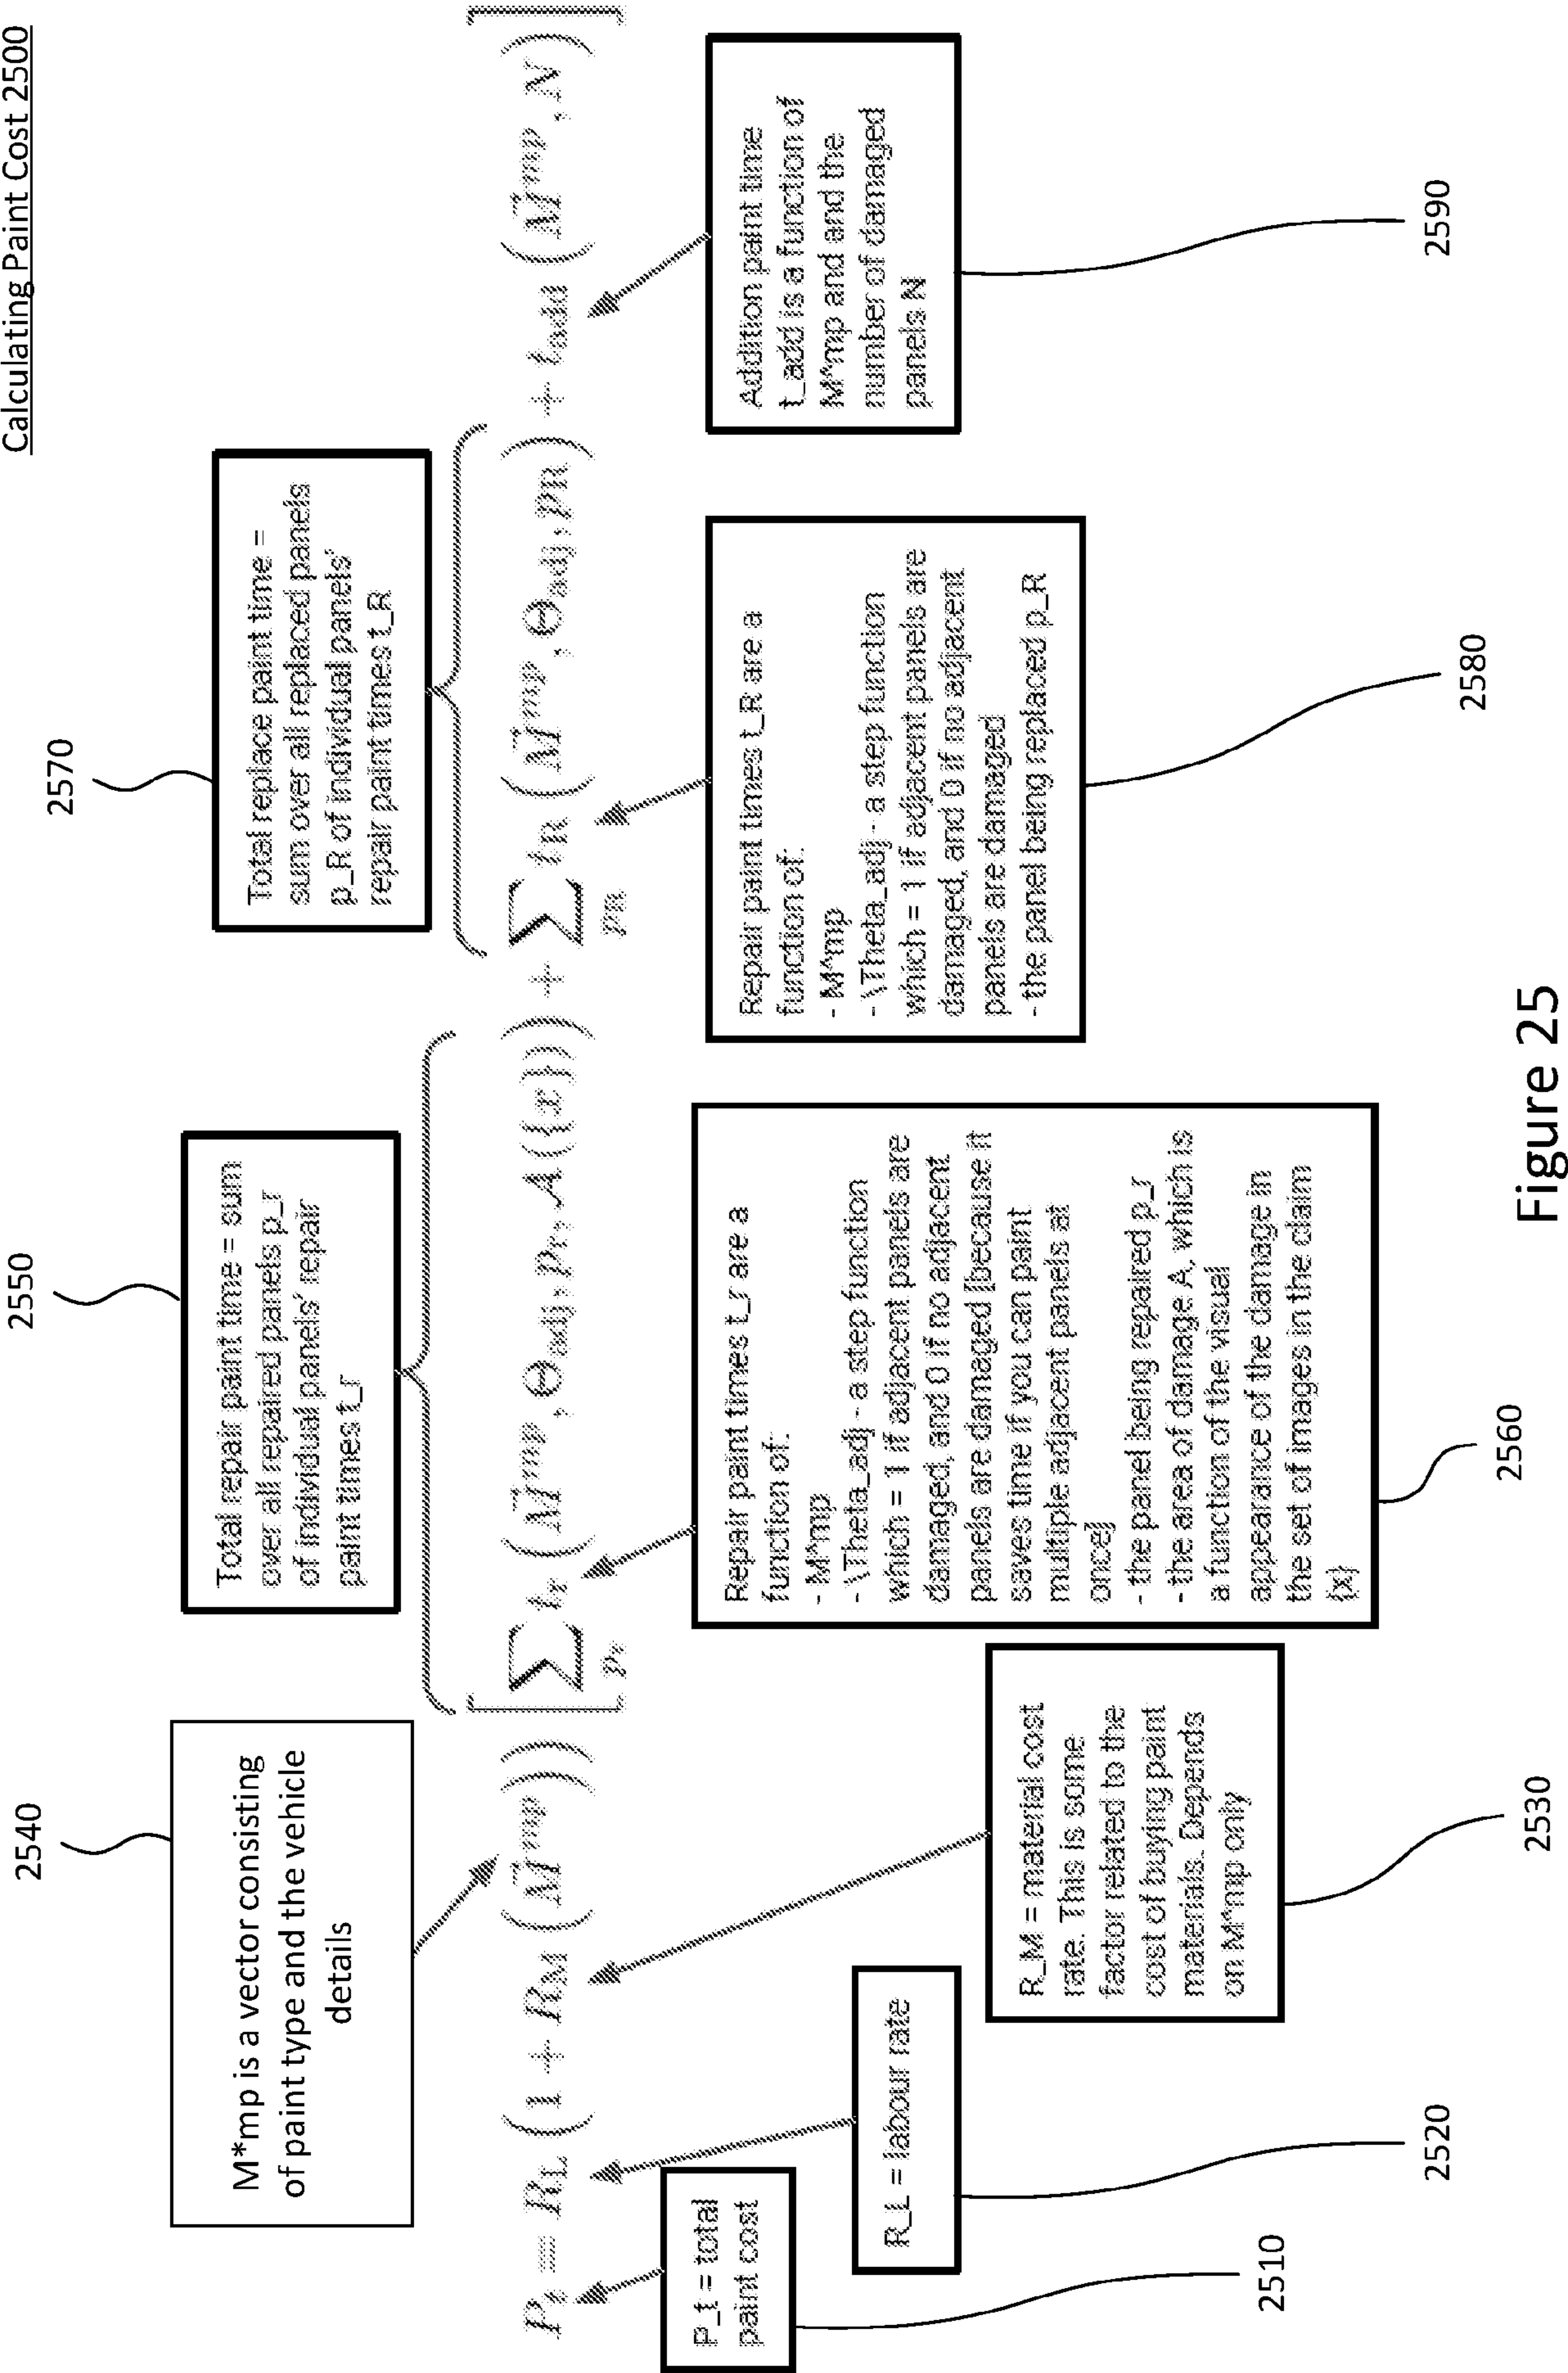


Figure 25

Overview of Damage
Determination Process 2600




Figure 26

Use of Cropping in Damage
Determination Process 2700




Figure 27

Use of Cropping/Segmentation in
Damage Determination Process
2800




Figure 28

Overview of Damage
Determination Process 2900

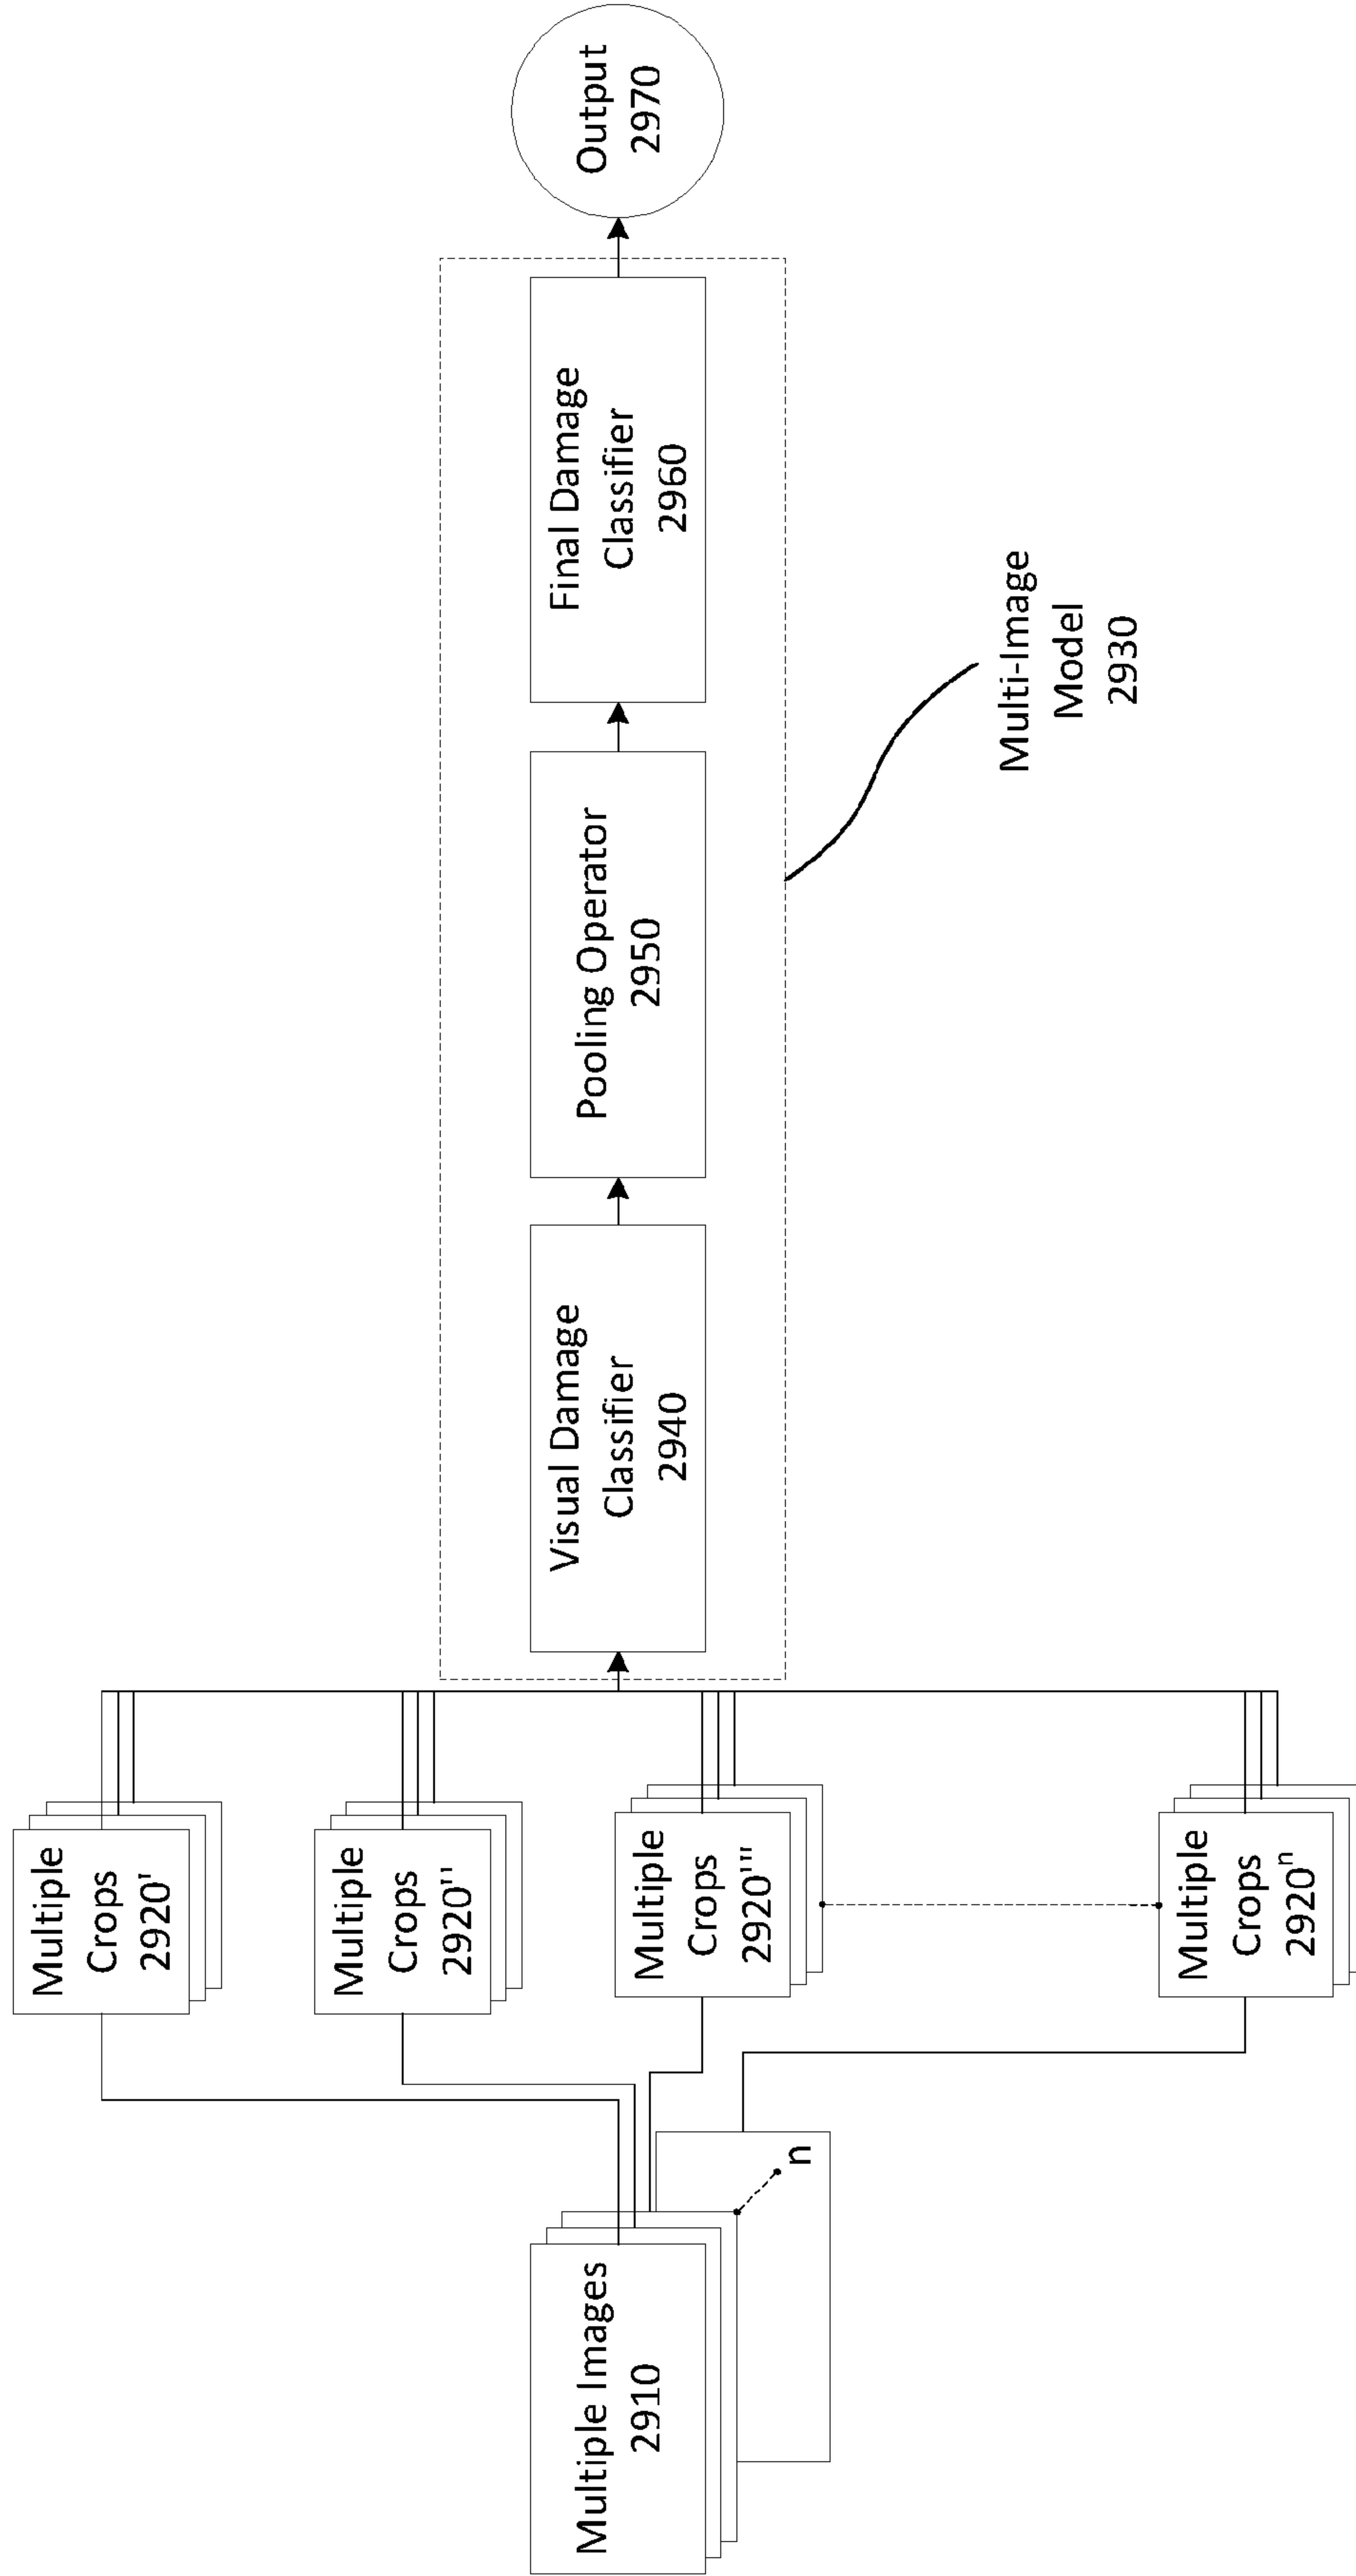


Figure 29

Overview of Damage
Determination Process 3000

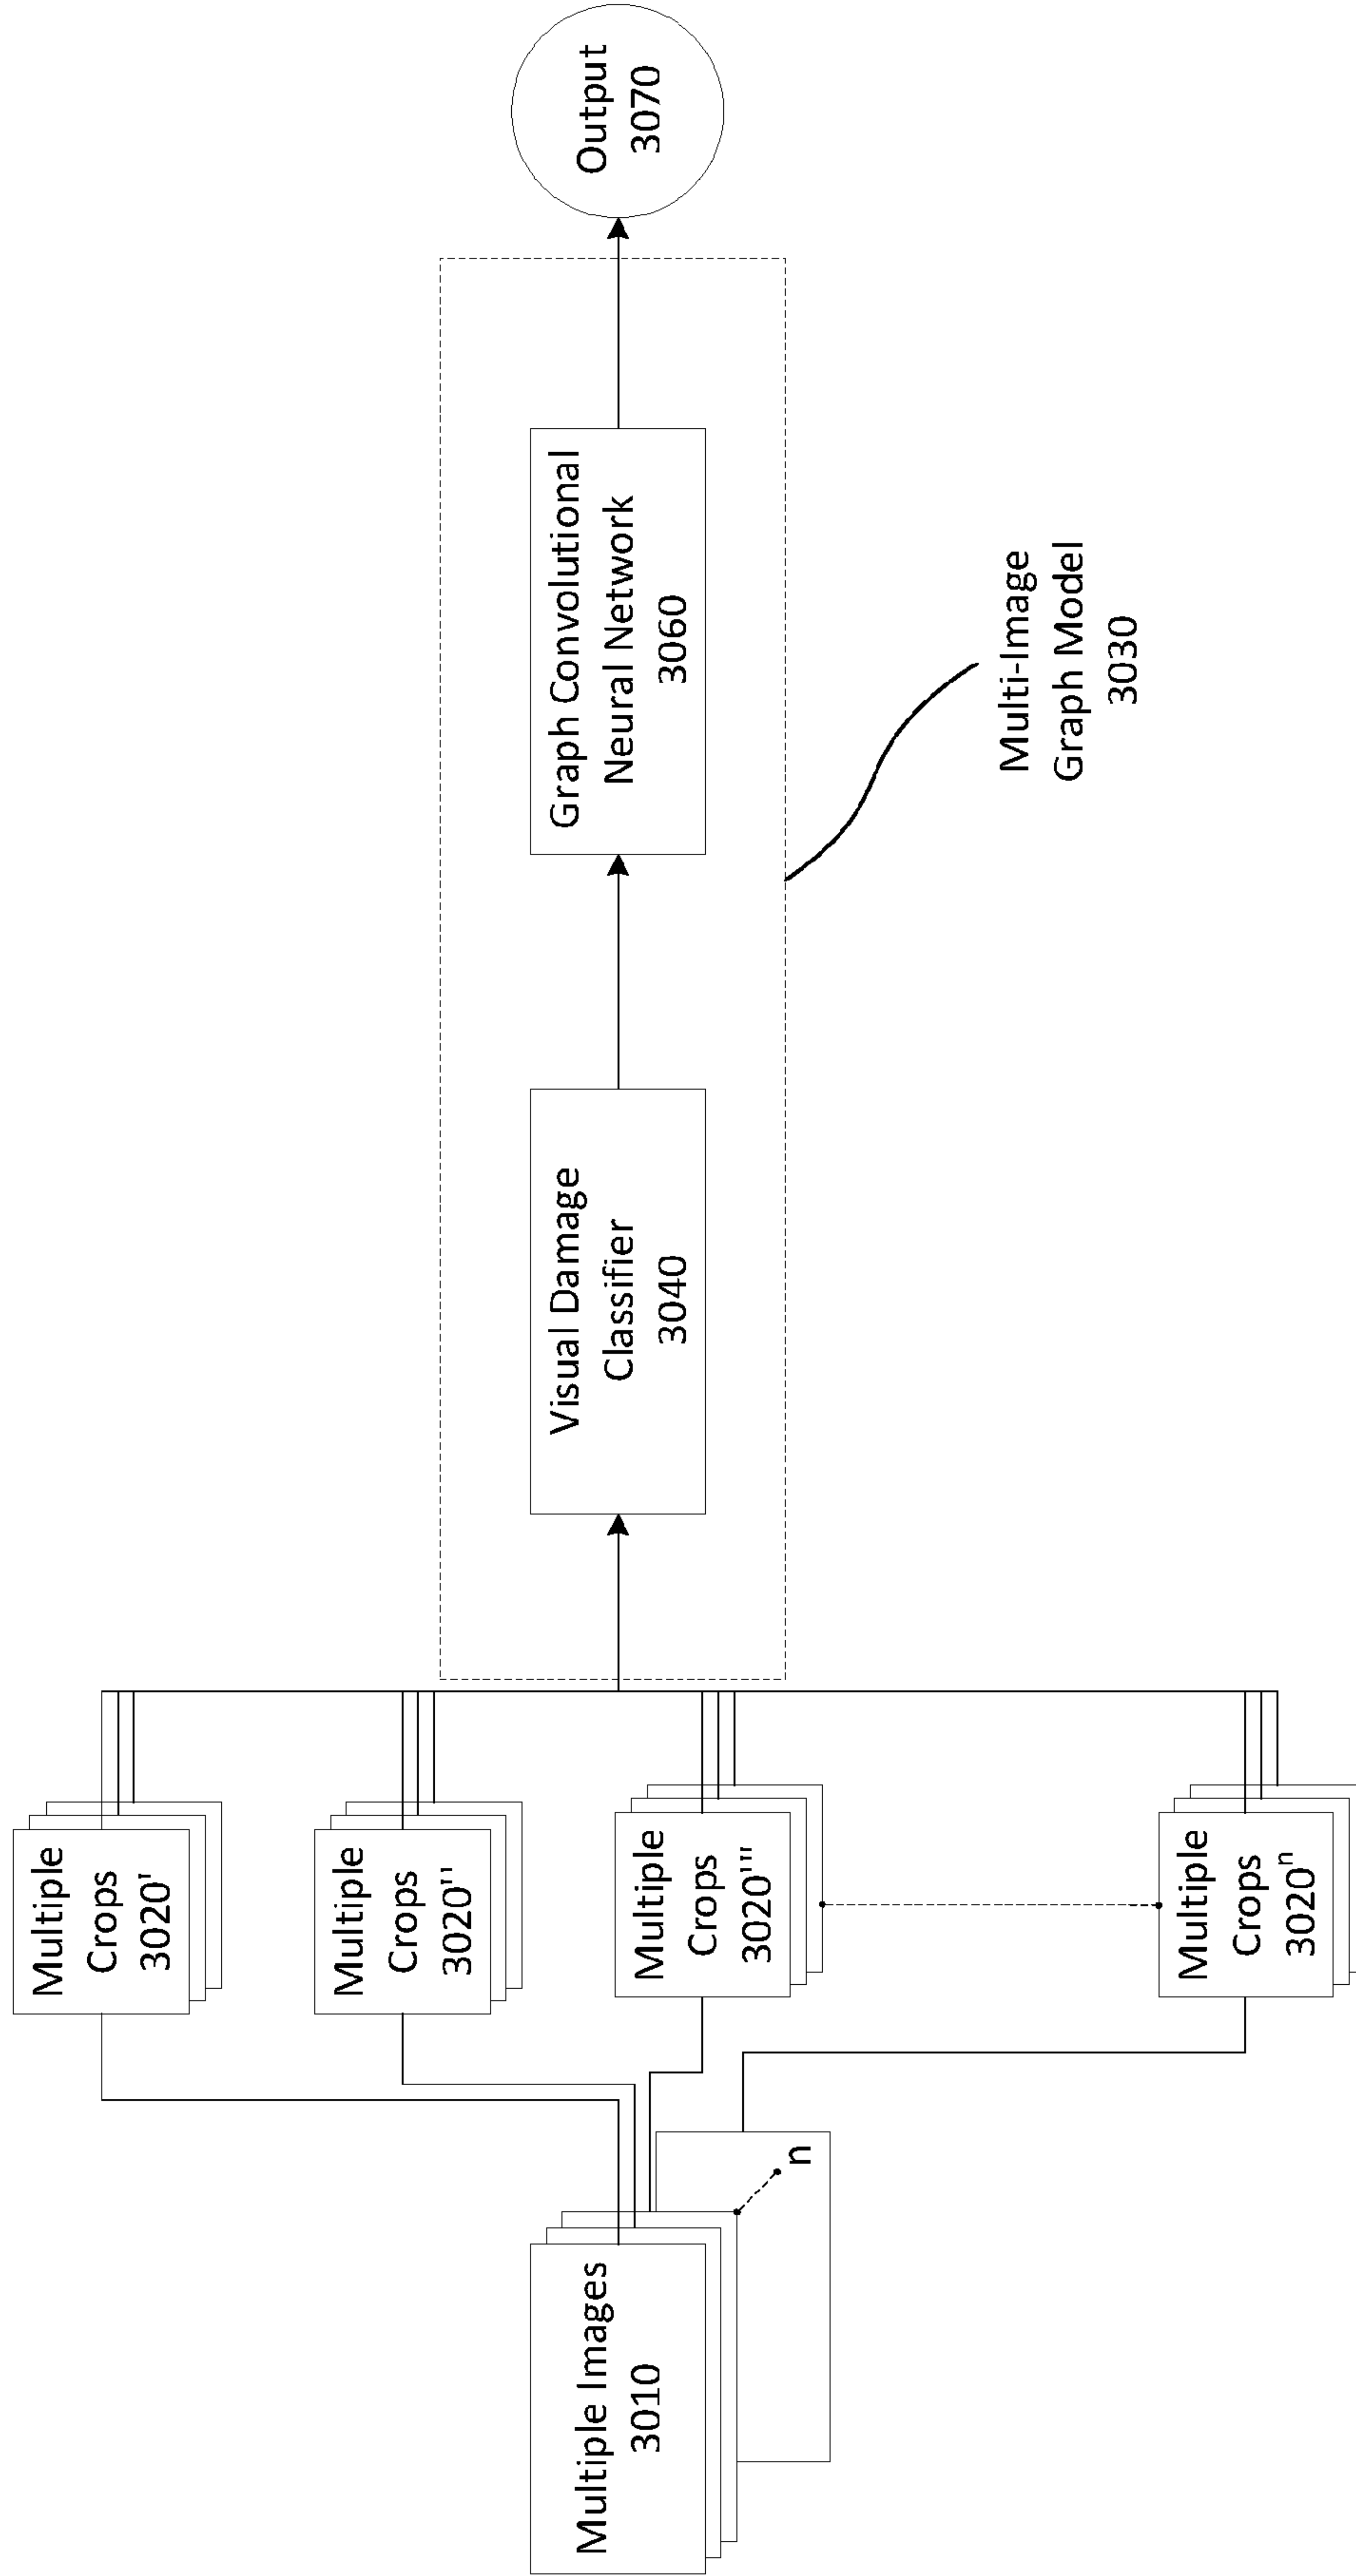


Figure 30

Using Language Model To Extract Information Relevant To Visual Damage 3100




Figure 31

Injecting Structured Accident Data
into Visual Damage Model 3200

32 / 37




Figure 32

Predicting damage to grouped
auxiliary parts associated with
normalised panel 3300

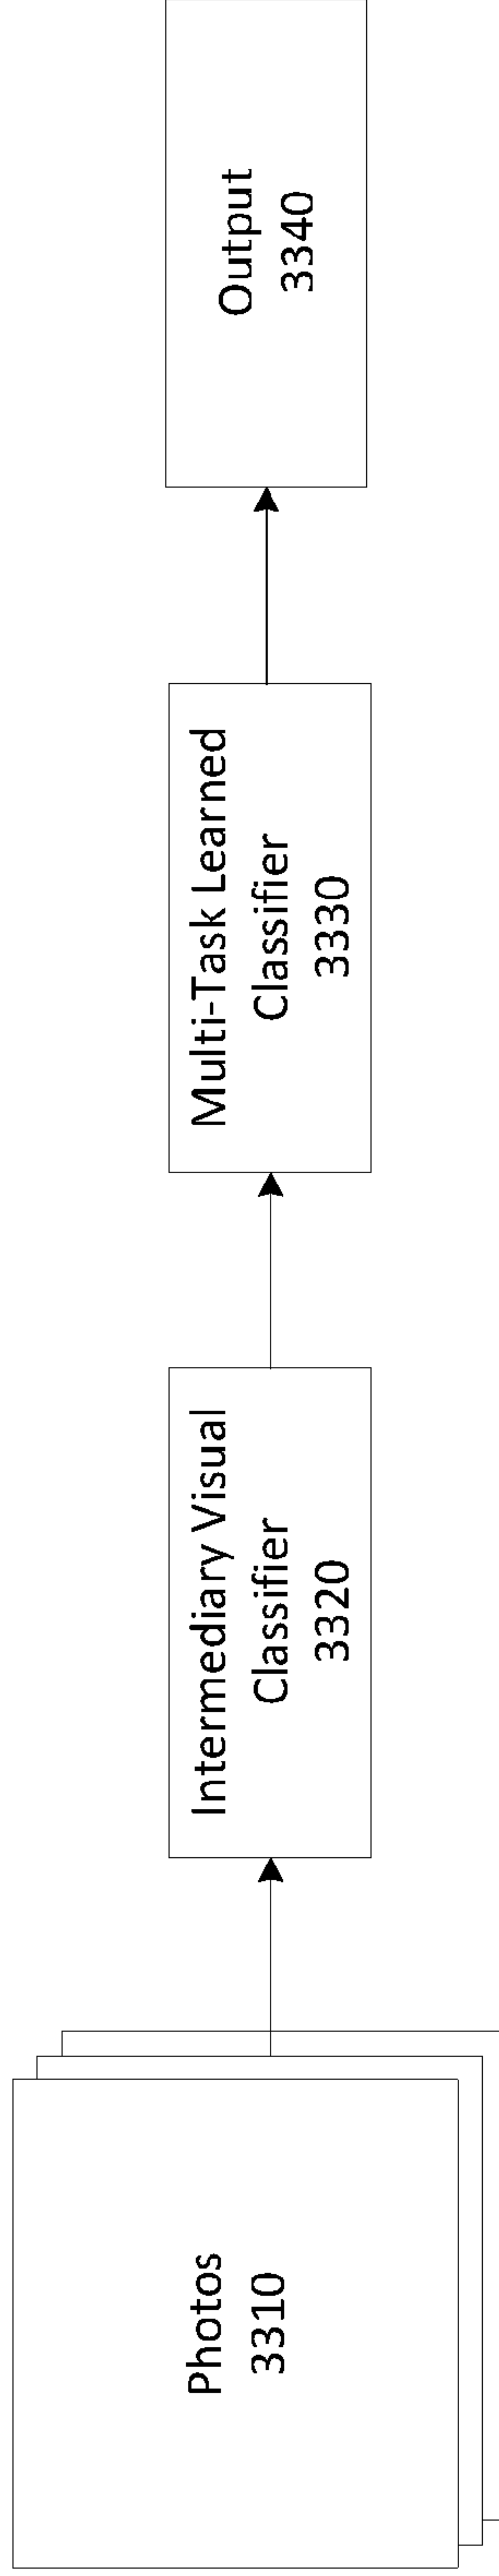


Figure 33

Example Multi-Task Architecture
3400


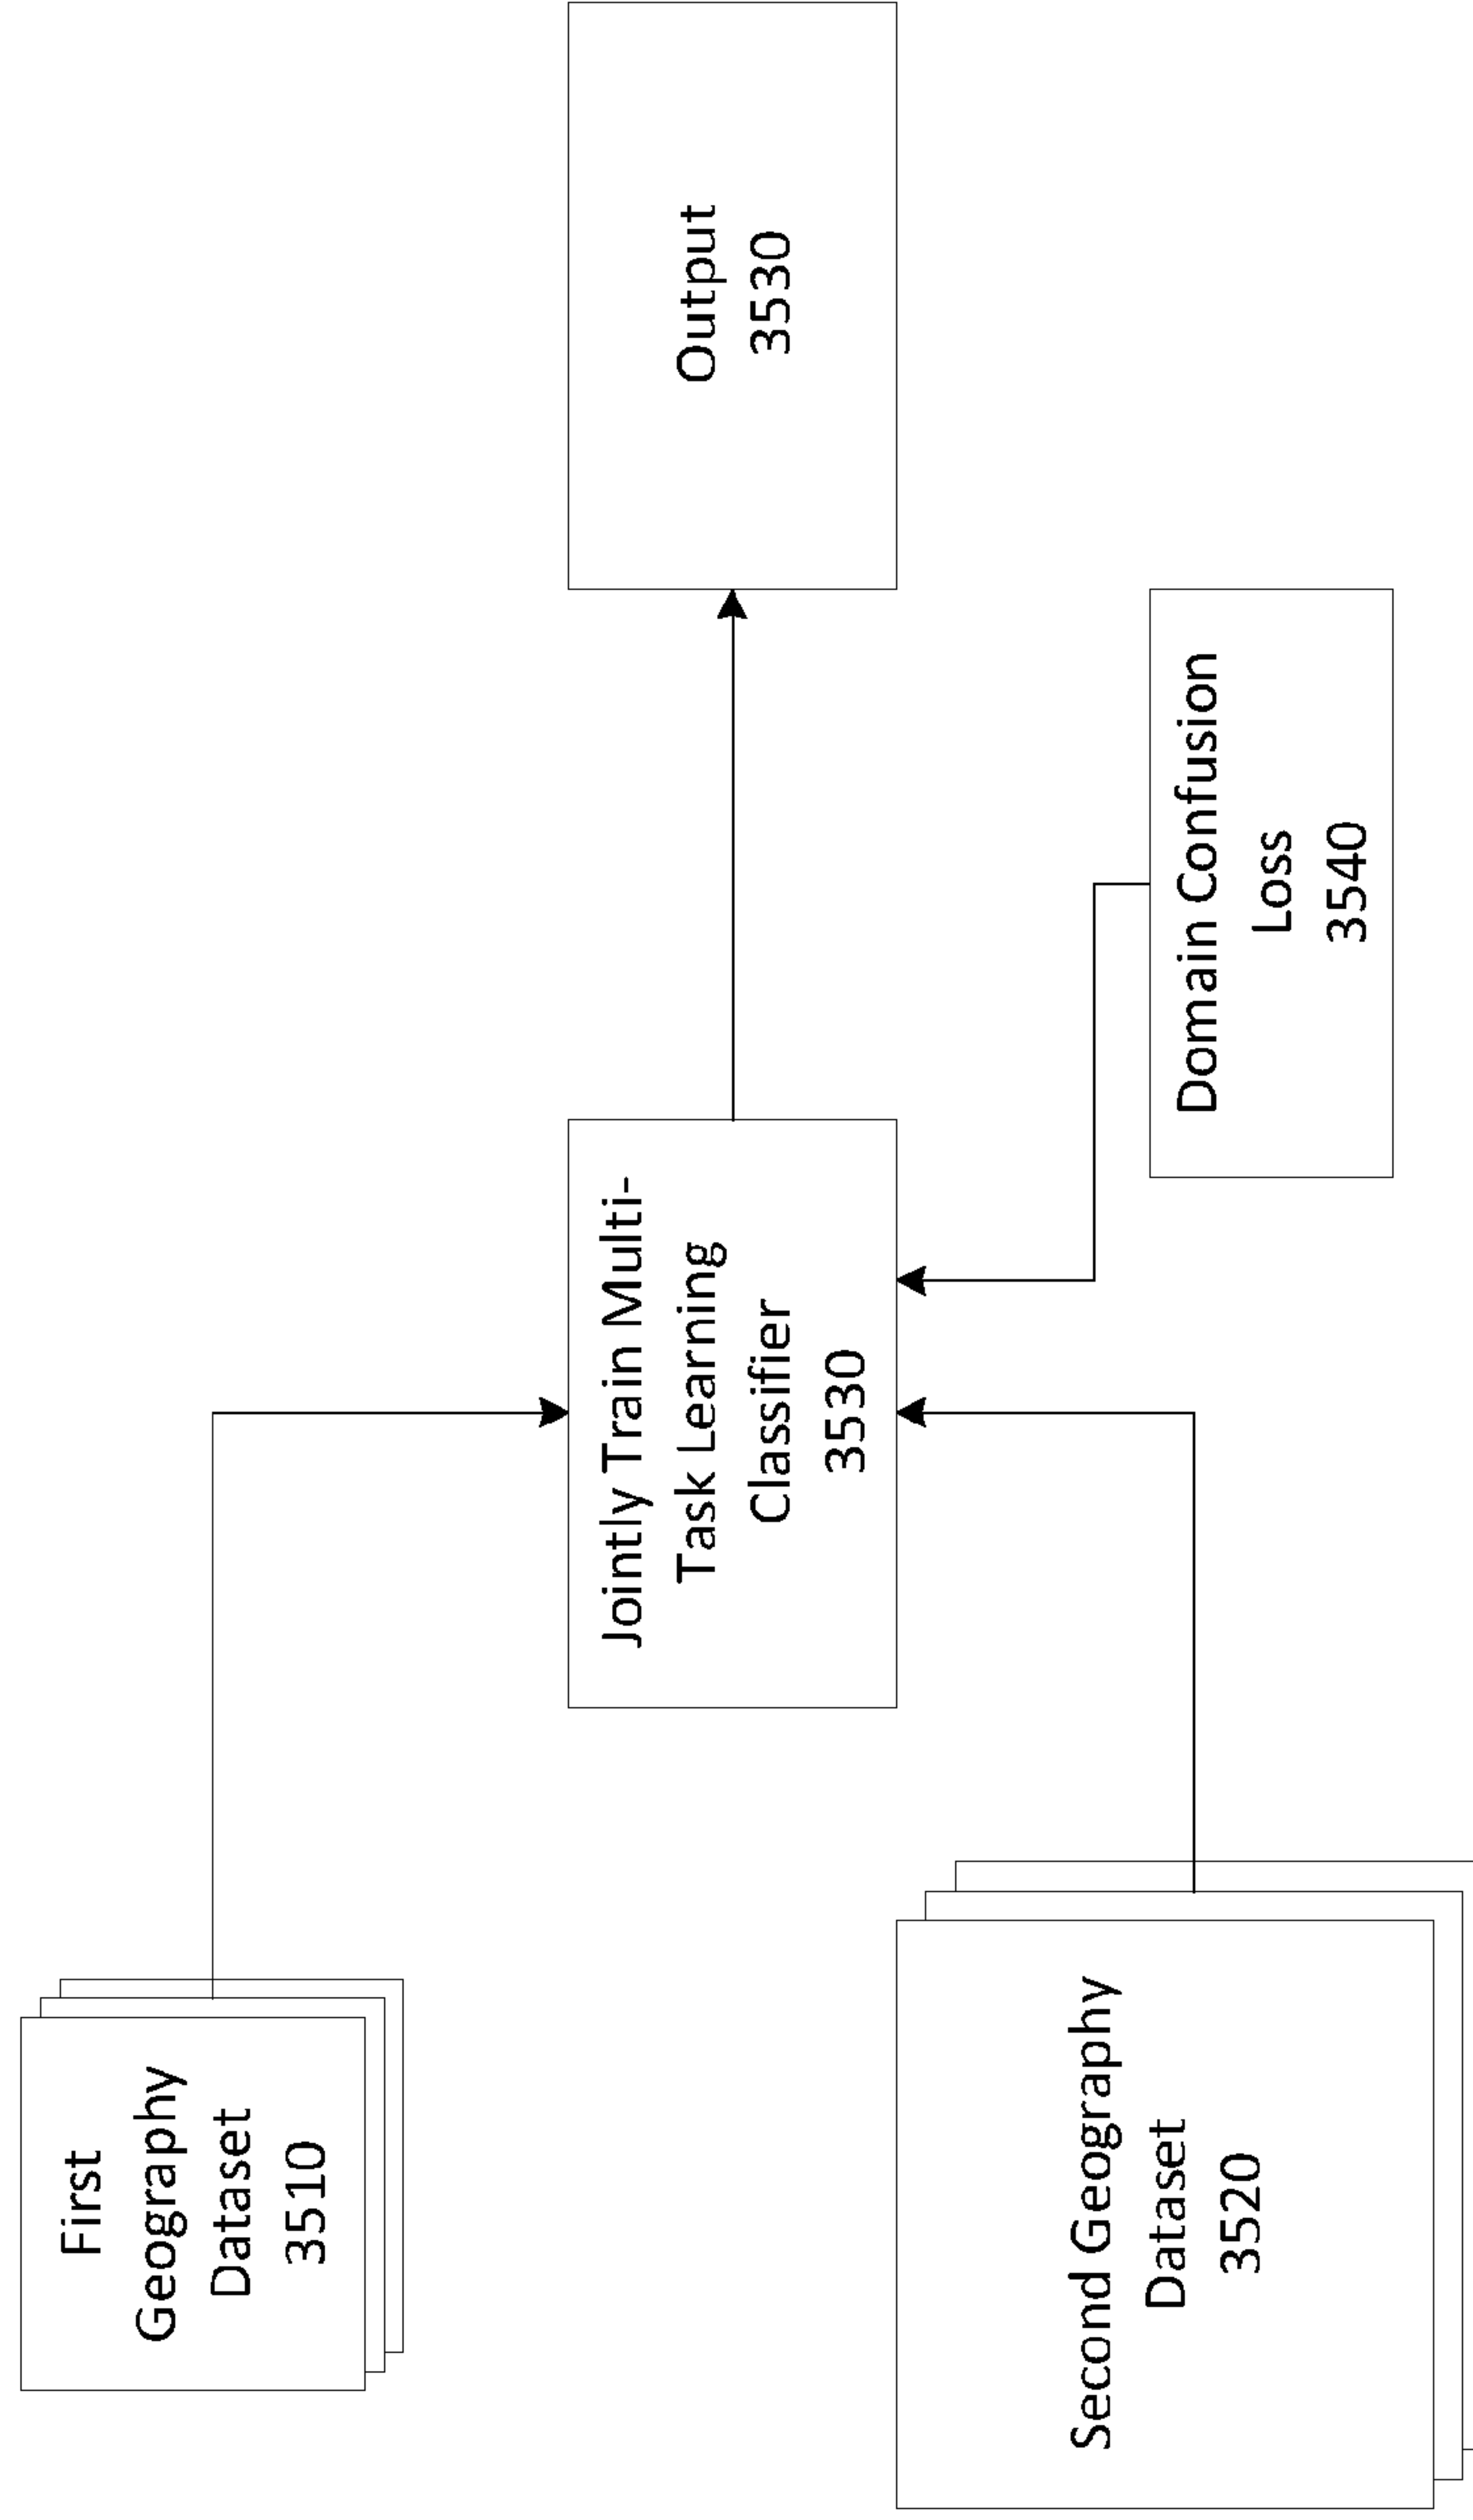


Figure 34

Jointly Training Across Variably Sized Datasets Using Domain Confusion Loss 3500

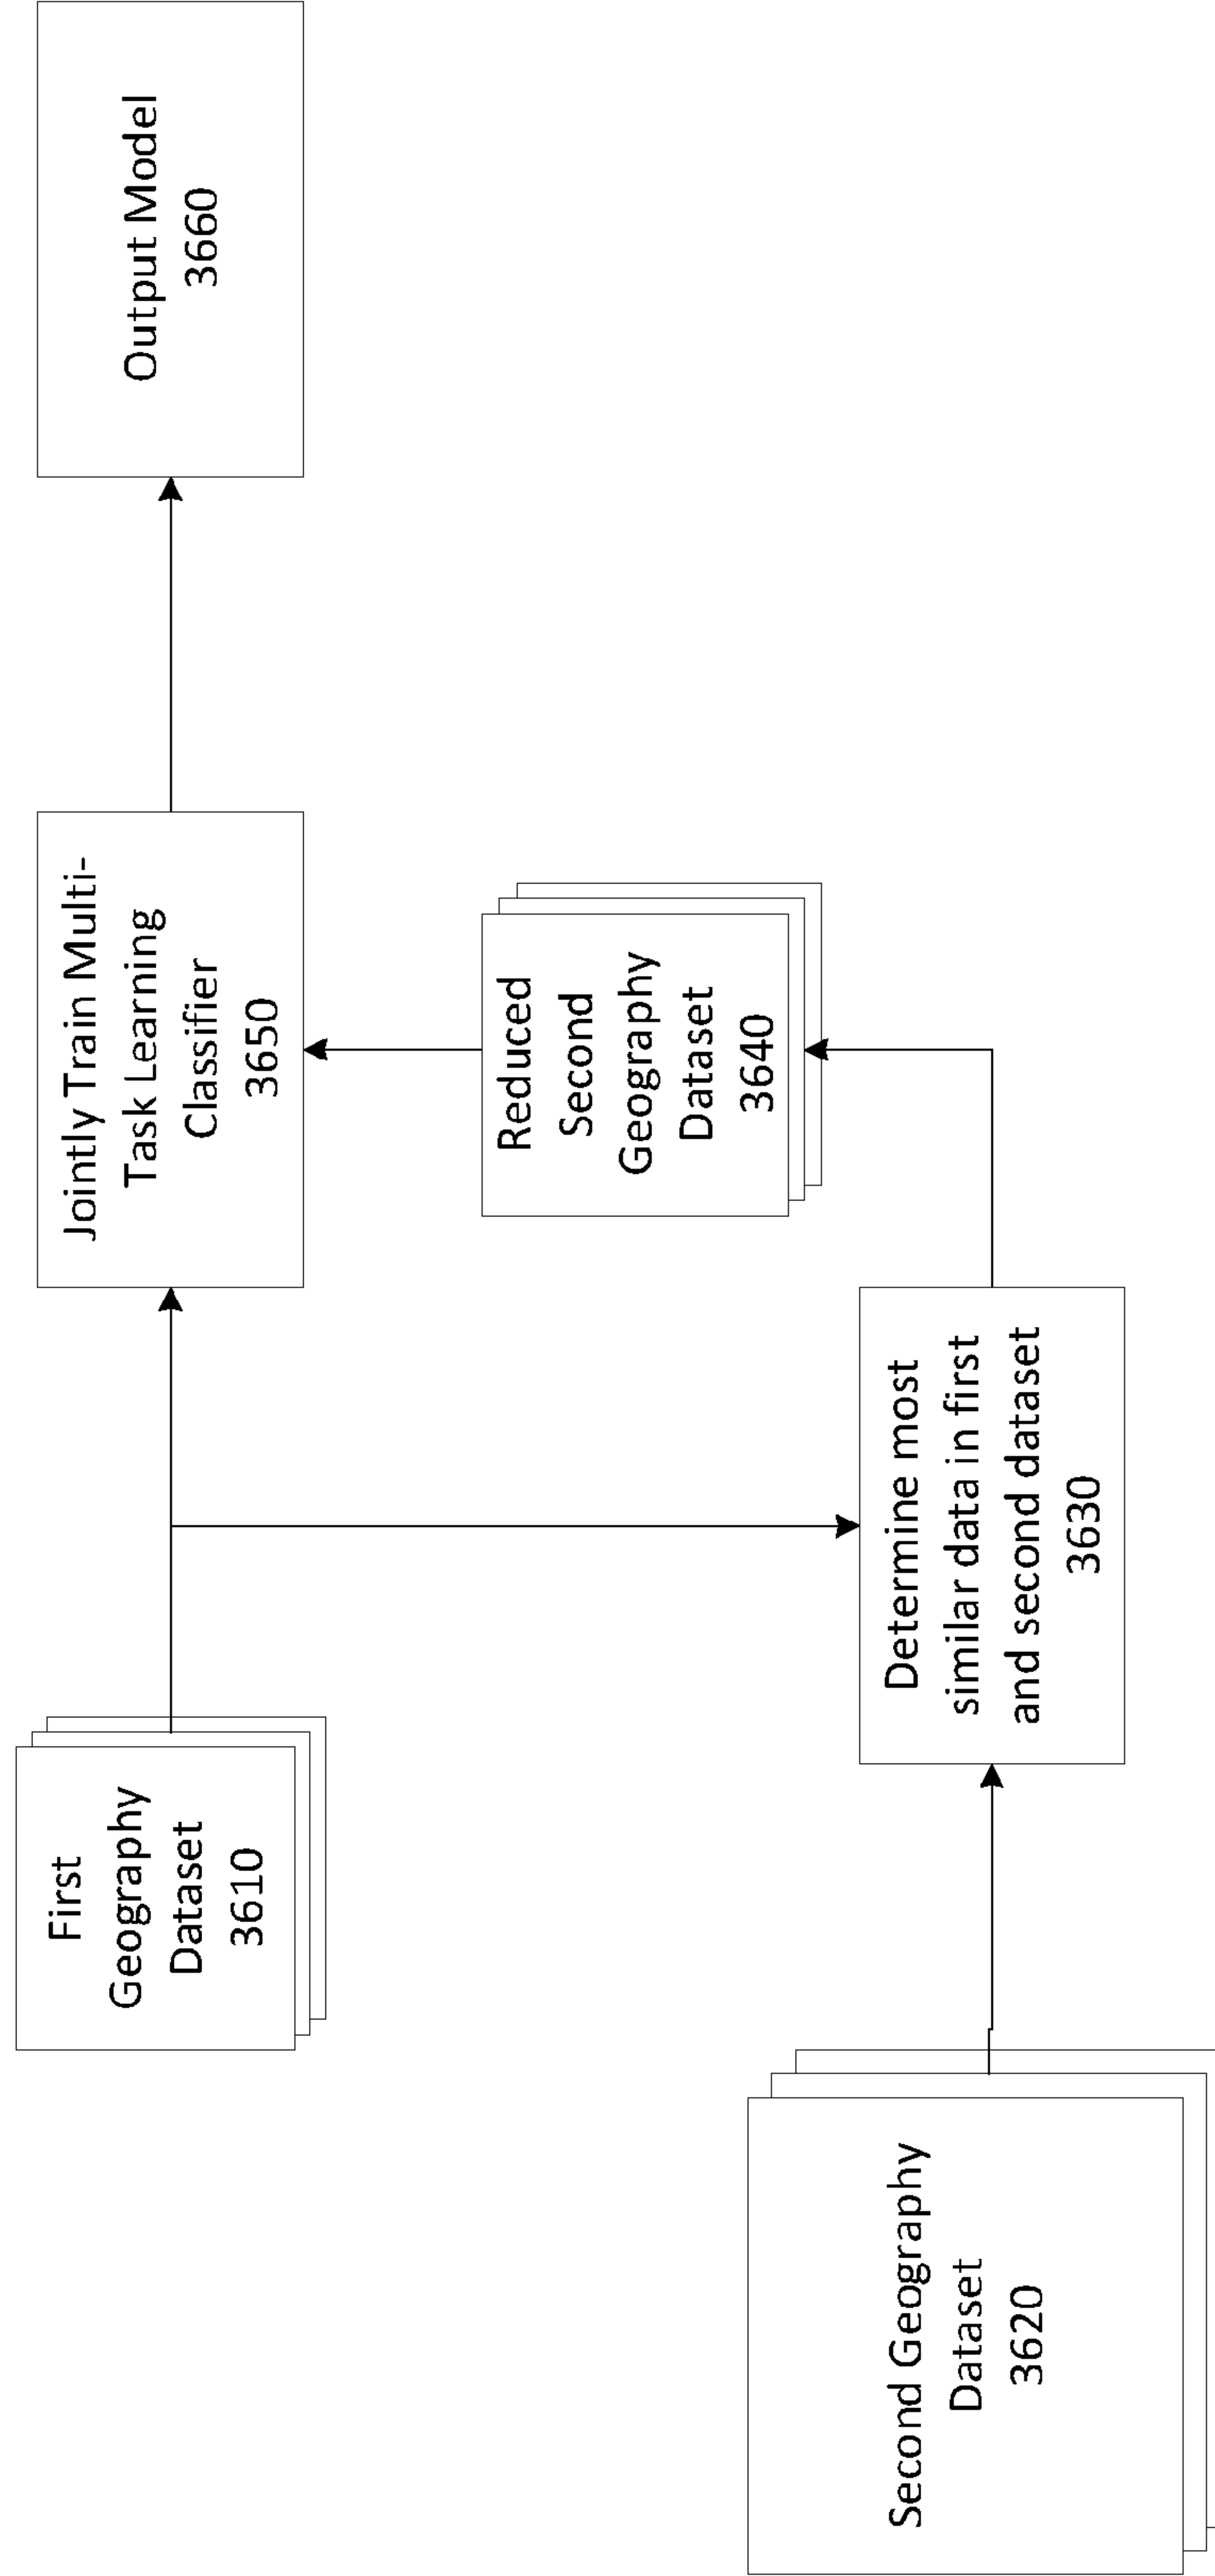
35 / 37



3560

Figure 35

Jointly Training Across Variably
Sized Datasets Using Most
Comparable Data 3600



3660

Figure 36

Example Domain Confusion Loss
Architecture 3700




Figure 37

INTERNATIONAL SEARCH REPORT

International application No PCT/GB2021/050010

A. CLASSIFICATION OF SUBJECT MATTER
 INV. G06Q10/00 G06Q40/08
 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)
 G06Q
 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, WPI Data

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 2018/293552 A1 (ZHANG HAITAO [CN] ET AL) 11 October 2018 (2018-10-11) figures 9-14 paragraphs [0005] - [0013] -----	1-16
X	US 2019/073641 A1 (UTKE JEAN [US]) 7 March 2019 (2019-03-07) figures 1A-3B,6,8 paragraphs [0017] - [0019] -----	1-16

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 25 March 2021	Date of mailing of the international search report 06/04/2021
--	--

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 Michalski, Stéphane
--	---

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No PCT/GB2021/050010

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2018293552	A1	11-10-2018	
		CN 107358596 A	17-11-2017
		EP 3520045 A1	07-08-2019
		JP 2020504358 A	06-02-2020
		JP 2020517015 A	11-06-2020
		KR 20190069457 A	19-06-2019
		PH 12019500971 A1	11-11-2019
		SG 11201903904V A	30-05-2019
		TW 201837786 A	16-10-2018
		US 2018293552 A1	11-10-2018
		US 2019213563 A1	11-07-2019
		WO 2018191435 A1	18-10-2018
		WO 2018191437 A1	18-10-2018

US 2019073641	A1	07-03-2019	
		EP 3676784 A1	08-07-2020
		US 2019073641 A1	07-03-2019
		WO 2019046400 A1	07-03-2019
