



US 20230306476A1

(19) **United States**

(12) **Patent Application Publication**
Osha

(10) **Pub. No.: US 2023/0306476 A1**

(43) **Pub. Date: Sep. 28, 2023**

(54) **SYSTEMS AND METHODS FOR VALUING AN ITEM**

(52) **U.S. Cl.**
CPC *G06Q 30/0278* (2013.01); *G06Q 30/0627* (2013.01); *G06Q 30/0643* (2013.01)

(71) Applicant: **Topmarq, Inc.**, Austin, TX (US)

(72) Inventor: **Quinn Osha**, Austin, TX (US)

(73) Assignee: **Topmarq, Inc.**, Austin, TX (US)

(21) Appl. No.: **18/180,672**

(22) Filed: **Mar. 8, 2023**

(57) **ABSTRACT**

Methods and systems are disclosed for valuing an item and for training a neural network to value an item. The systems may include a client web interface configured to input identification information for the item and a digital representation of the item; and a computer system configured to: obtain item specifications for the item using the identification information; determine a preliminary condition indication based on the item specifications; determine, from the digital representation using a machine learning (ML) engine, a first item attribute; determine, from the first item attribute, a first condition indication adjustment; determine an adjusted condition indication based on the preliminary condition indication and the first condition indication adjustment; obtain market valuation data of the item based on the identification information; and determine a value of the item based on the market valuation data and the adjusted condition indication.

Related U.S. Application Data

(60) Provisional application No. 63/317,945, filed on Mar. 8, 2022.

Publication Classification

(51) **Int. Cl.**
G06Q 30/02 (2006.01)
G06Q 30/0601 (2006.01)

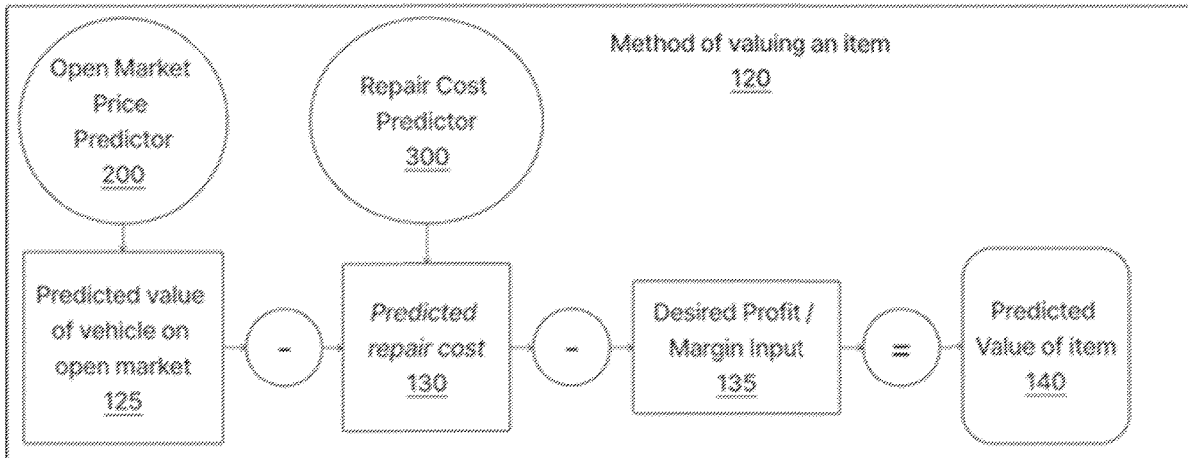
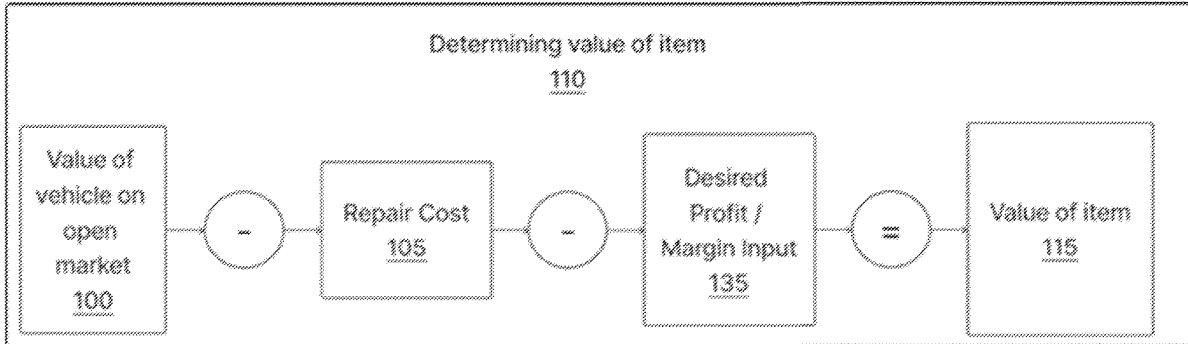


FIG. 1A





FIG. 1B



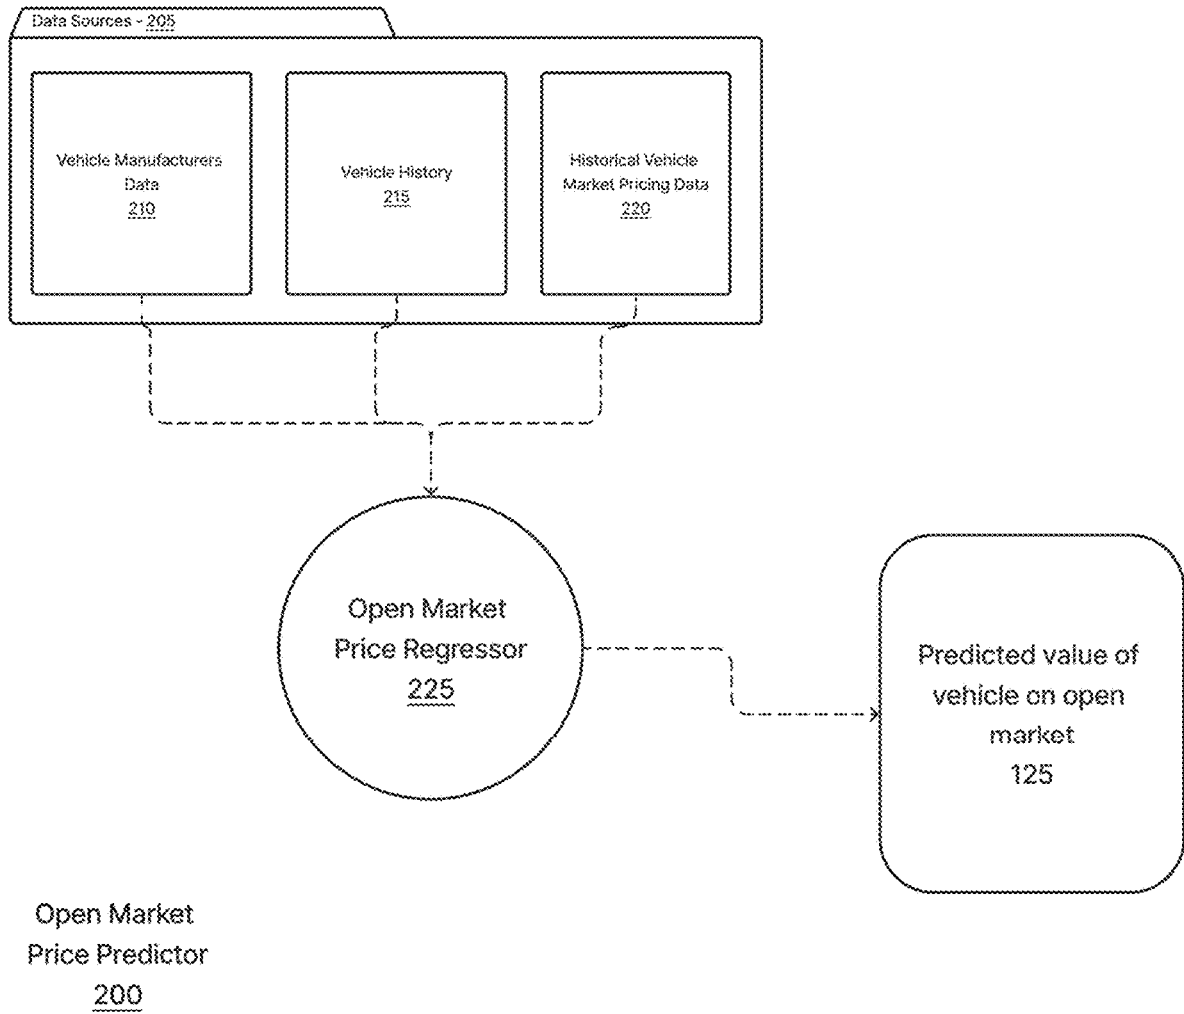


FIG. 2



Repair Cost
Predictor
300

FIG. 3

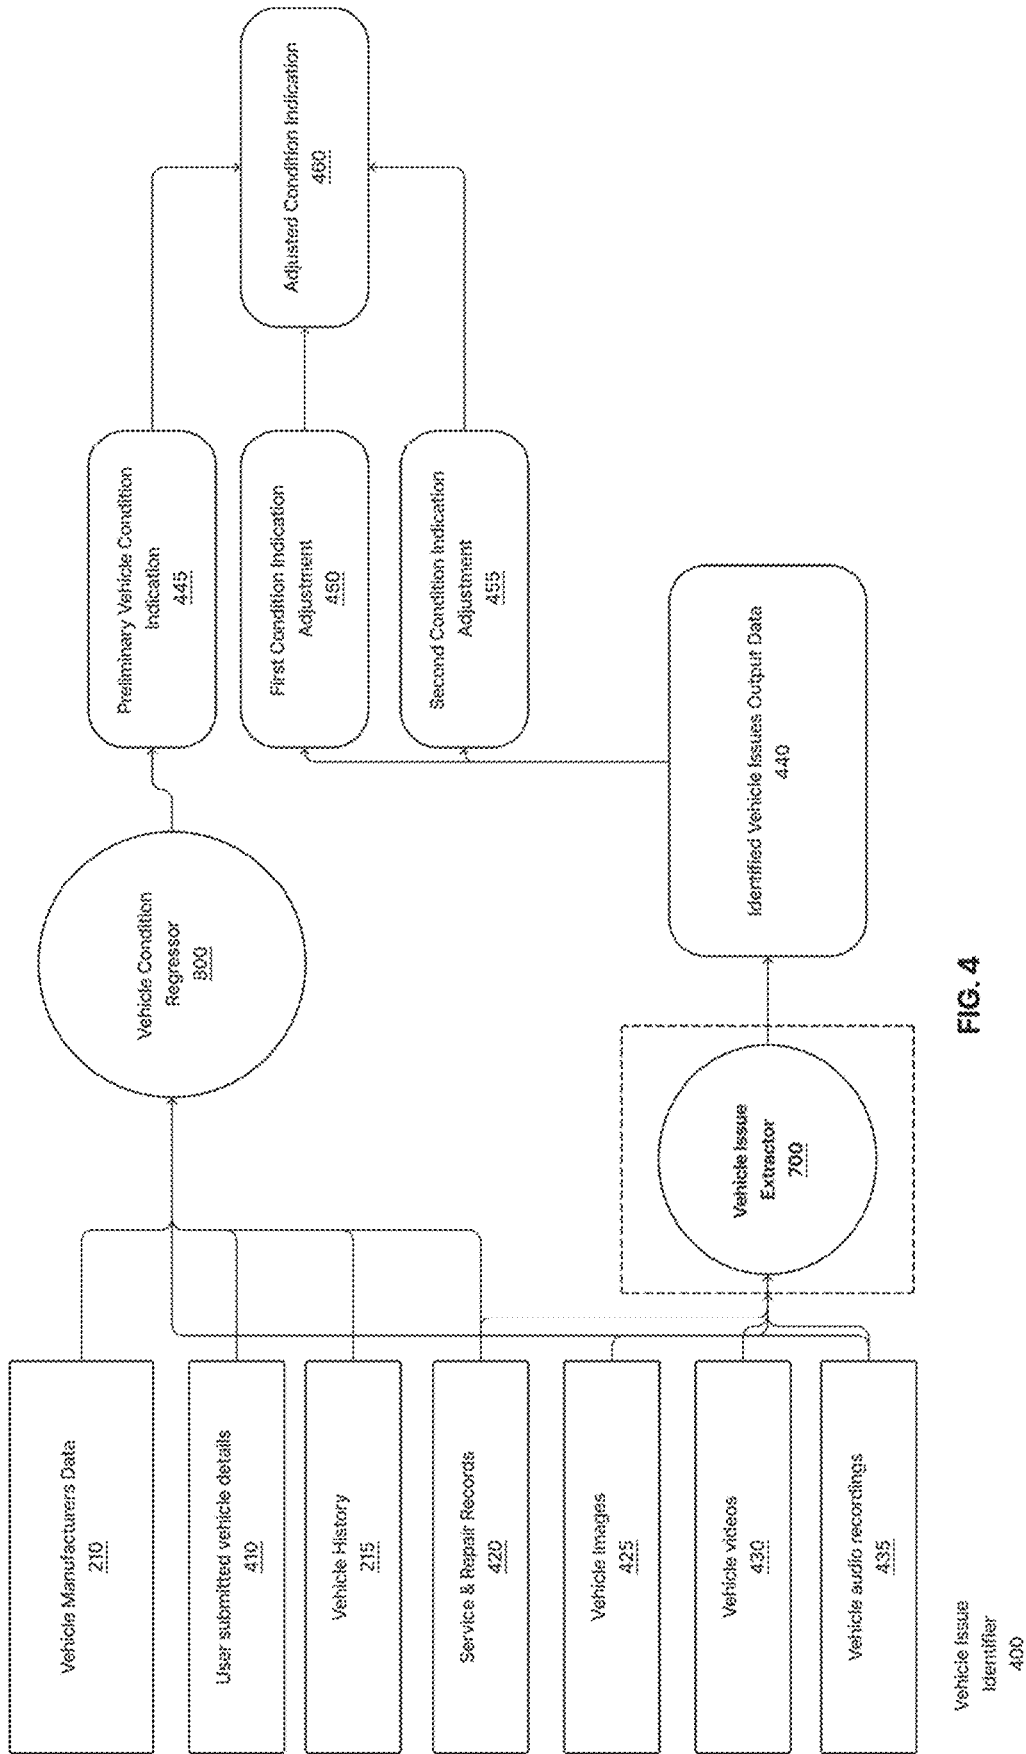
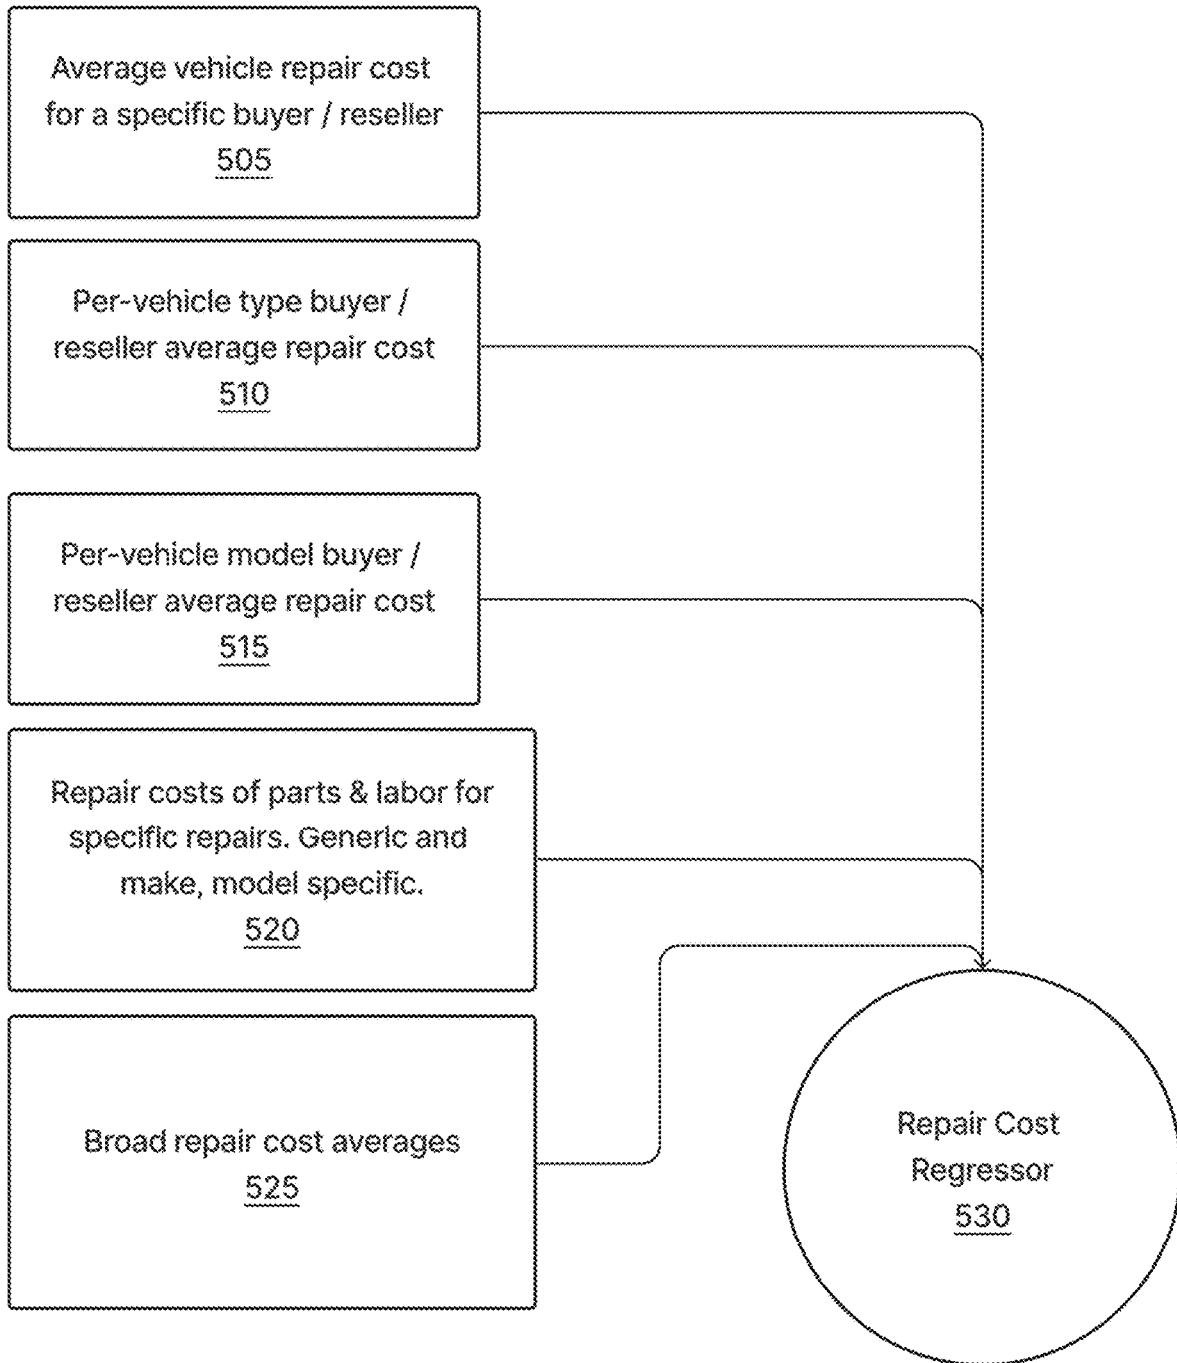



FIG. 4



Issue Repair Cost
Predictor
500

FIG. 5



Predictor Operational
Mode
60

FIG. 6

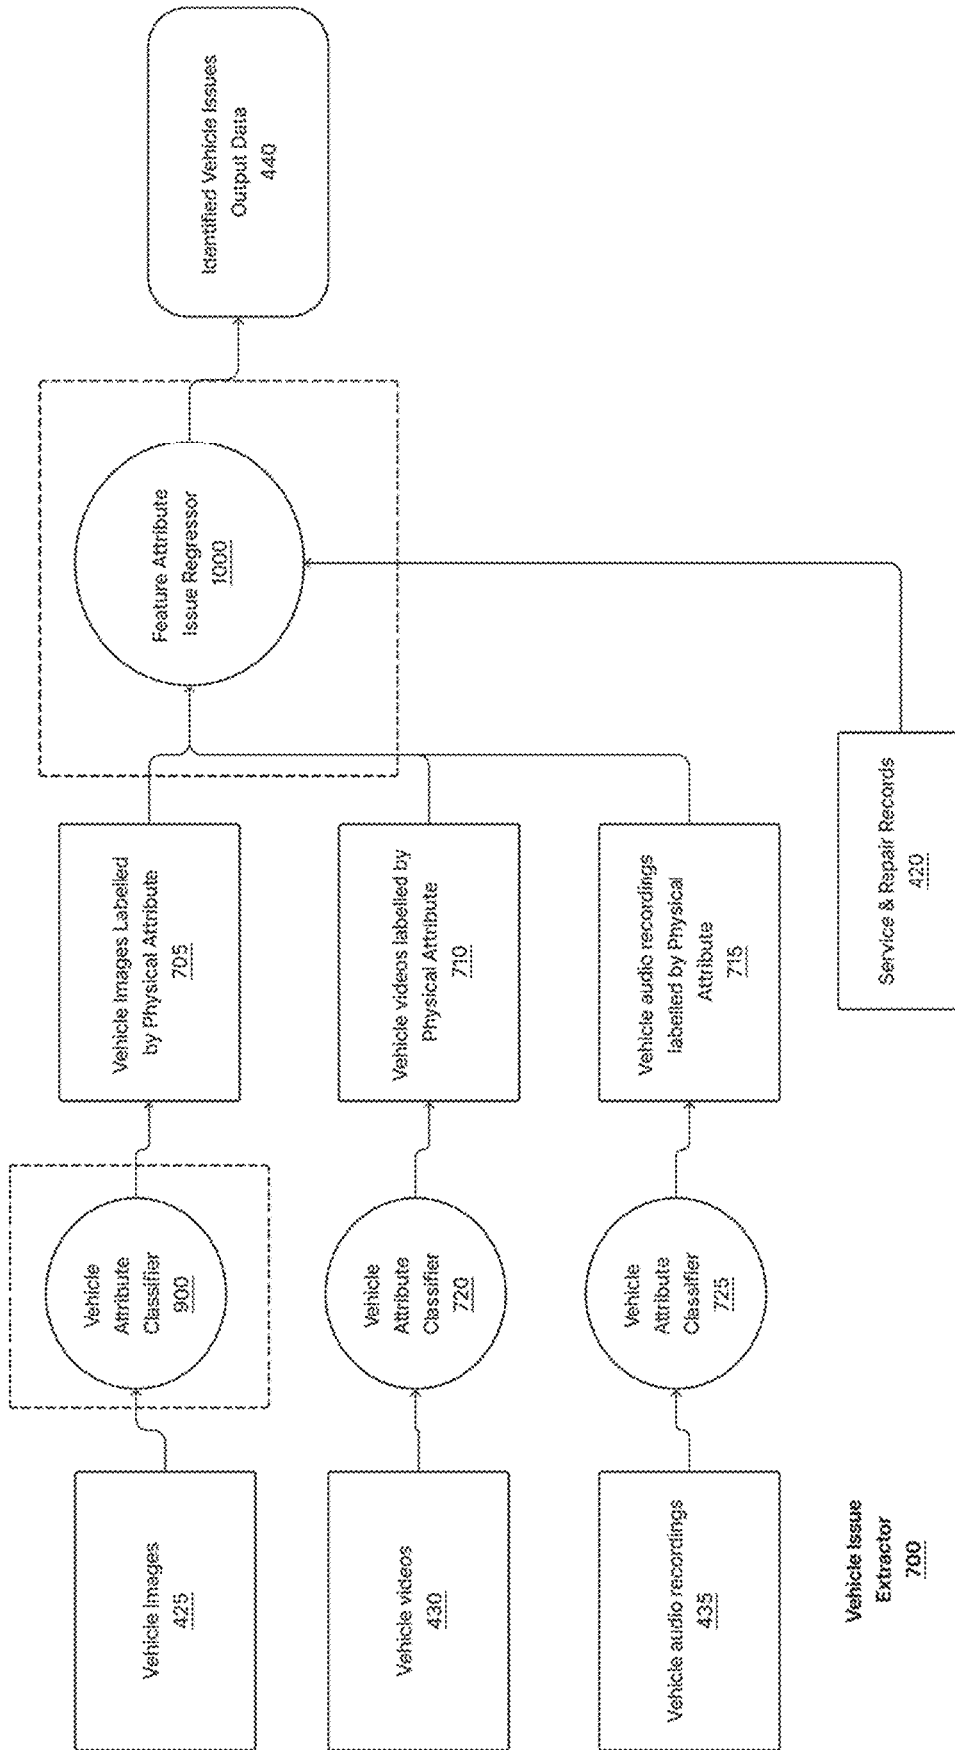


FIG. 7

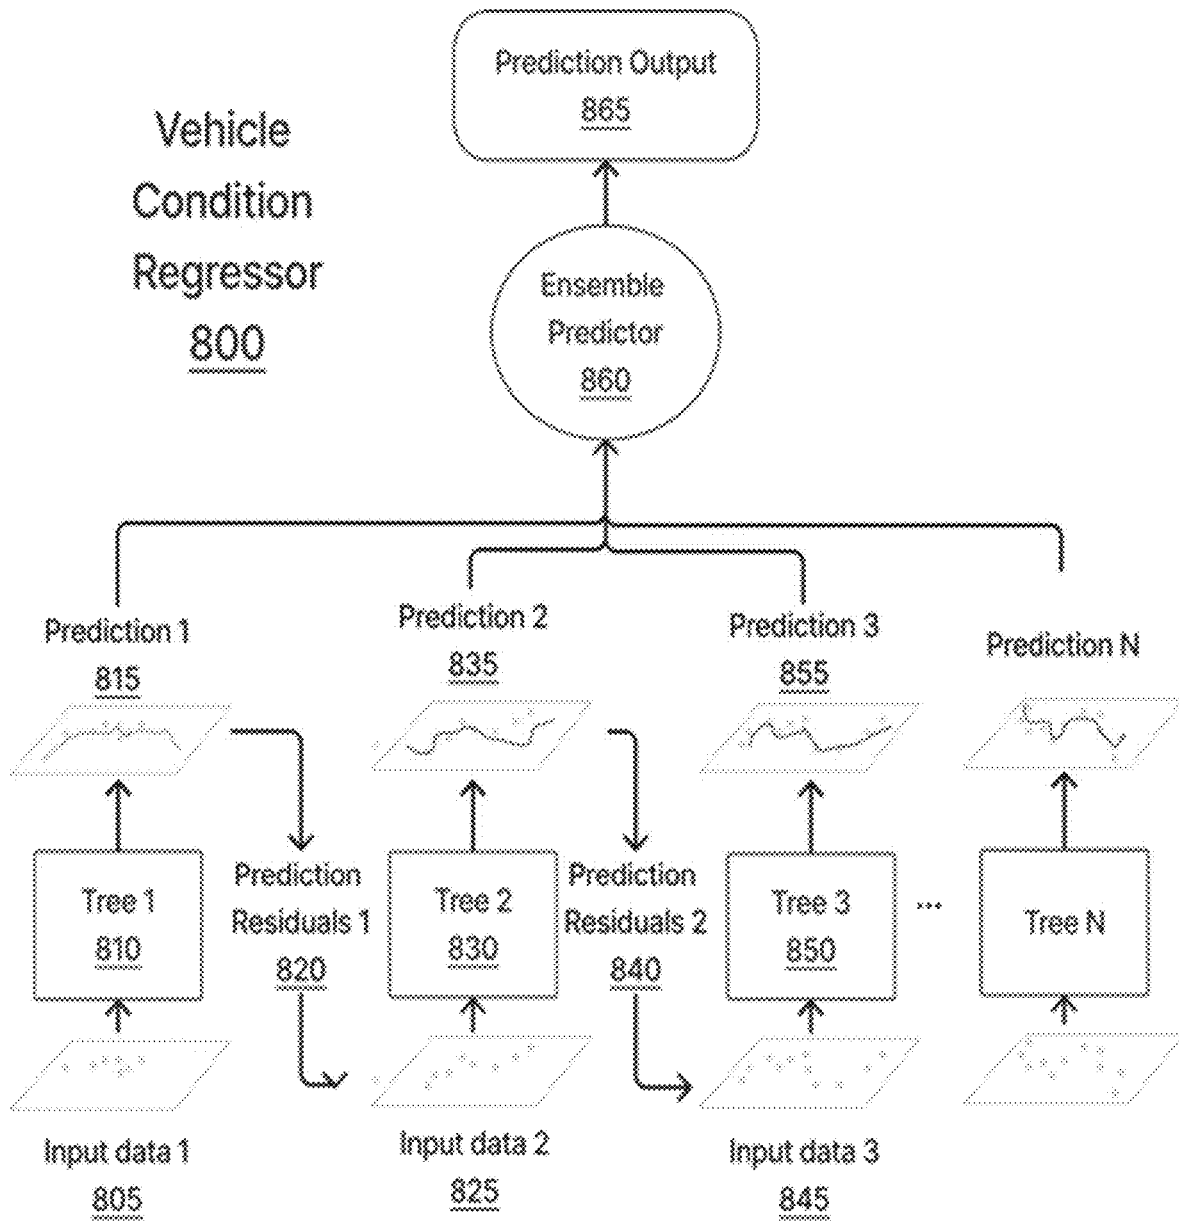


FIG. 8

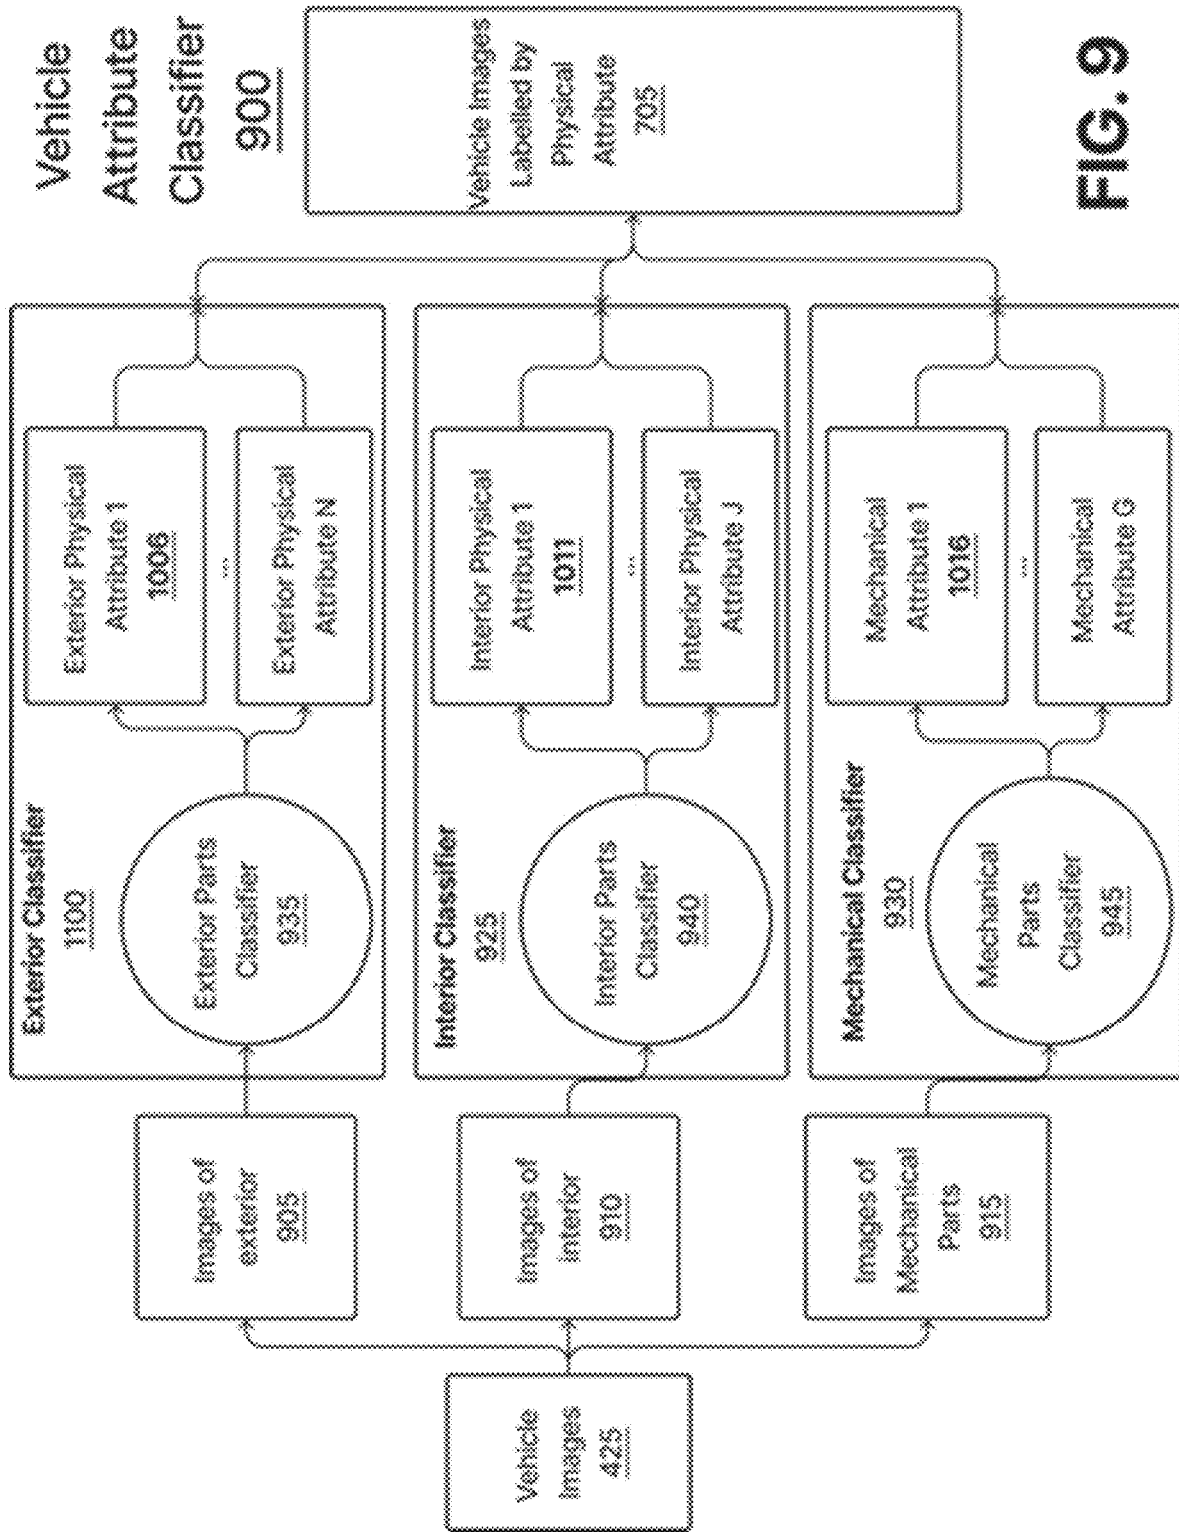


FIG. 9

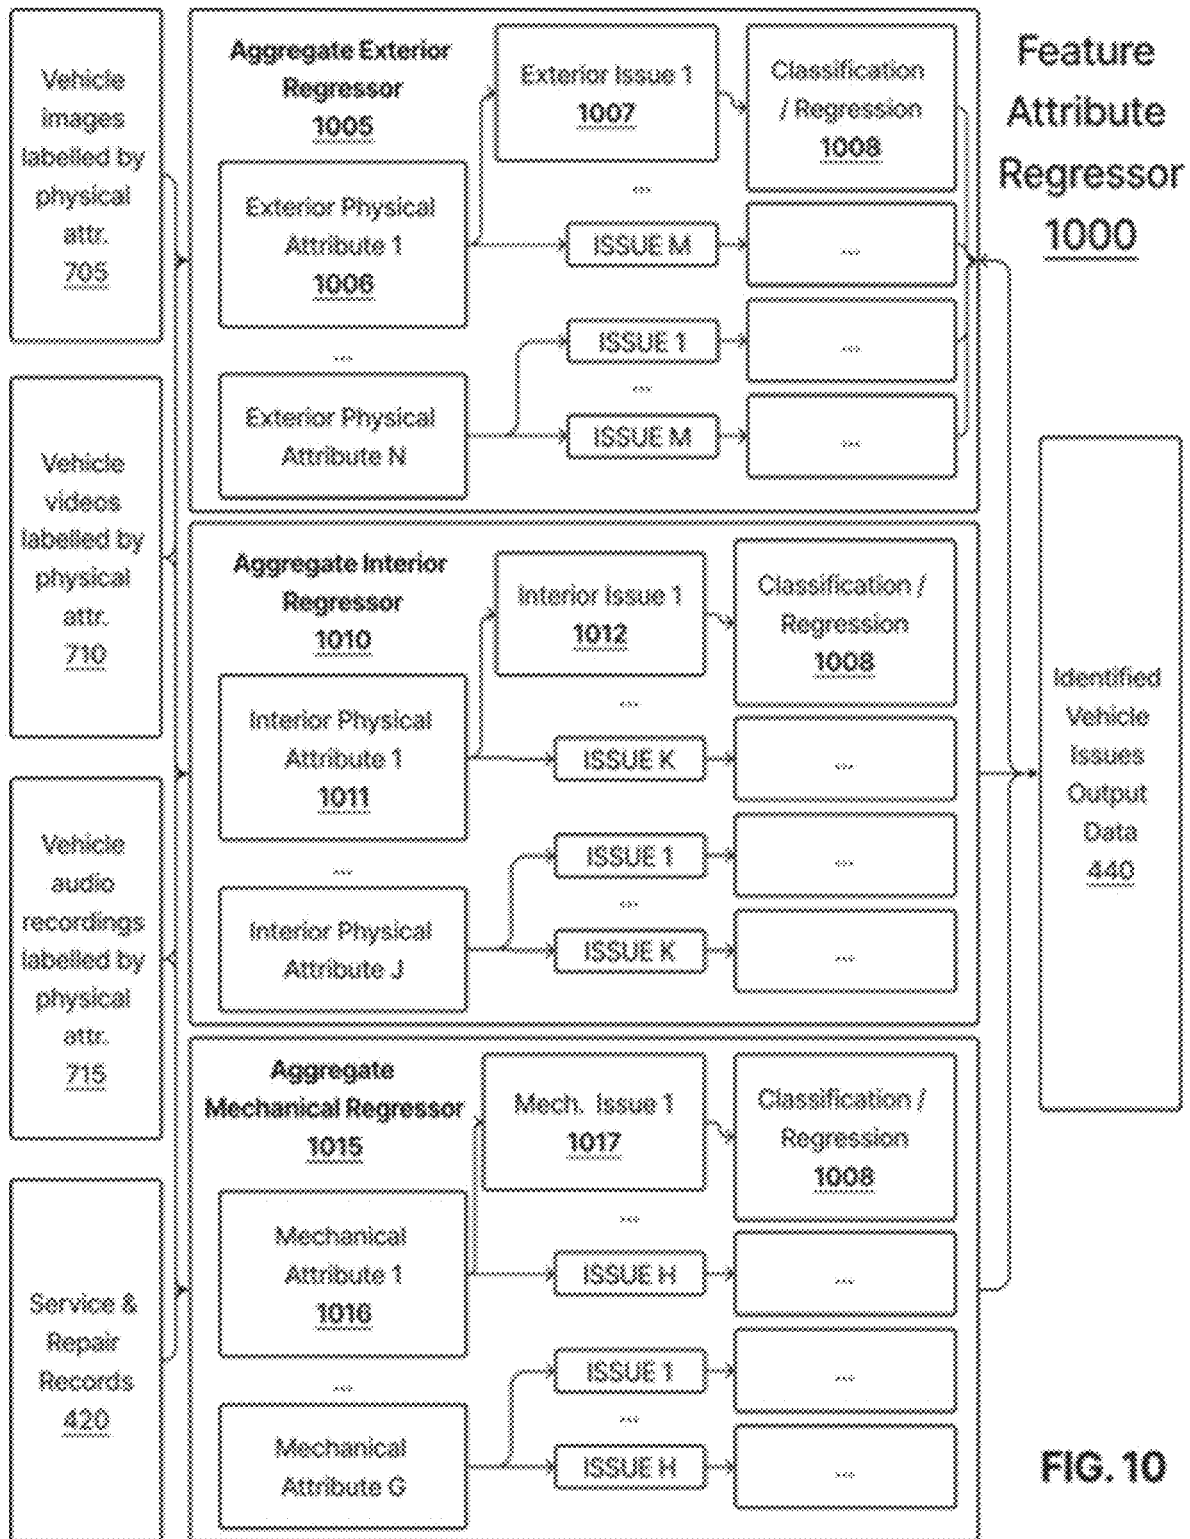



FIG. 10



Exterior Classifier 1100

FIG. 11




FIG. 12


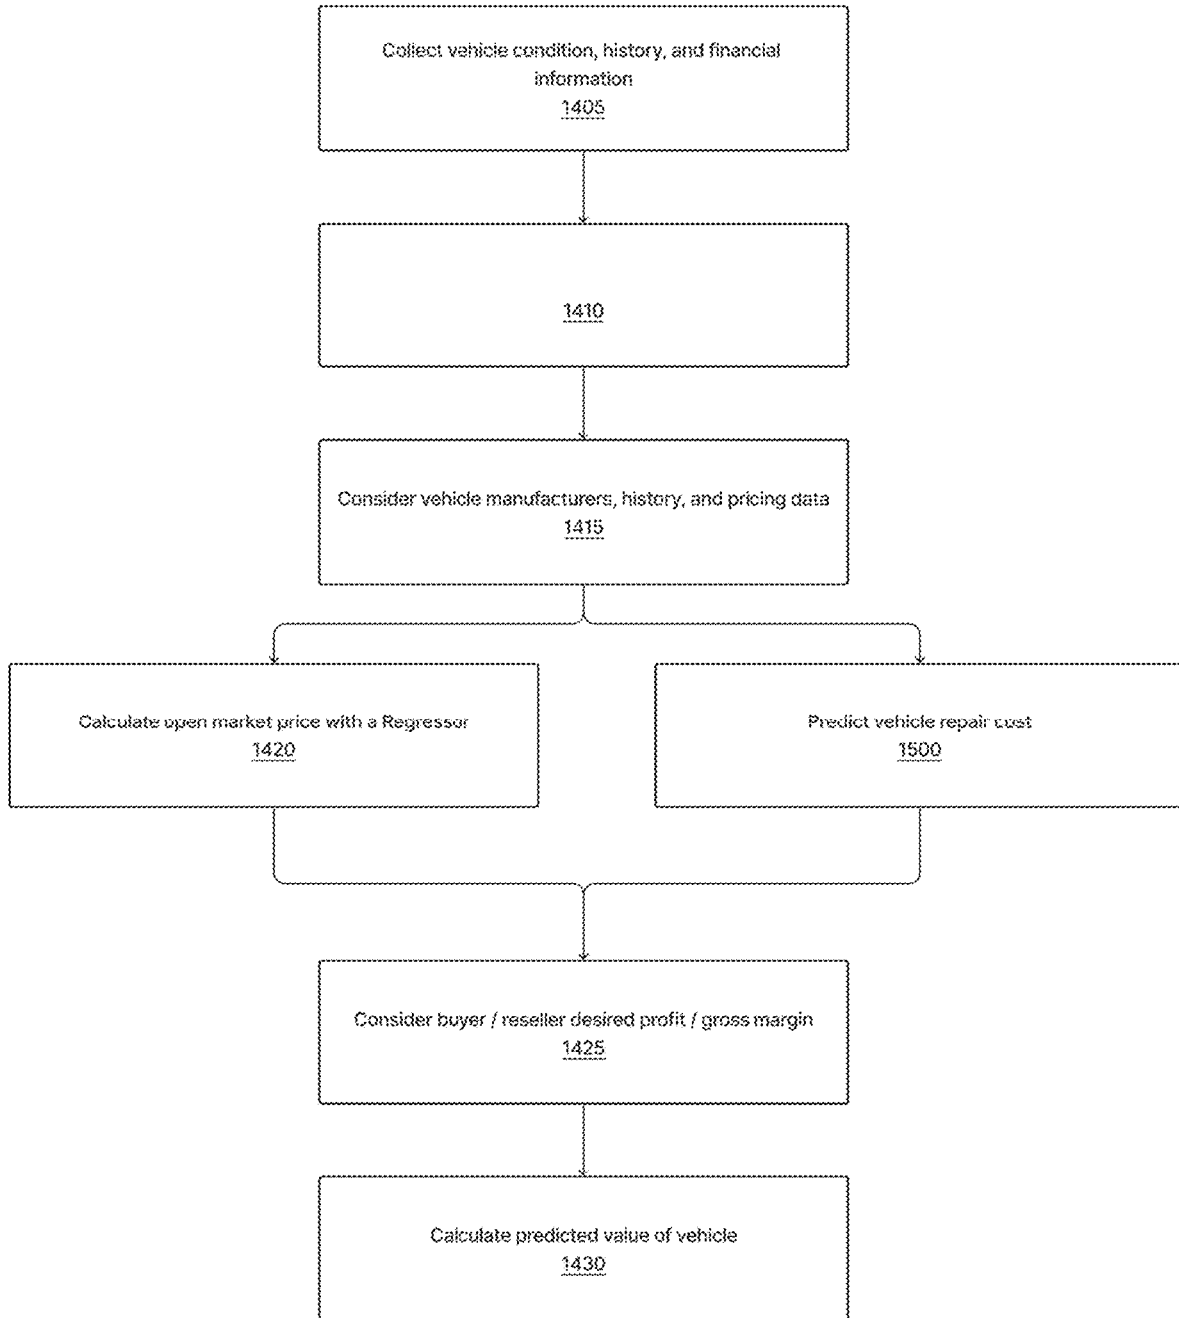


FIG. 13



1400

FIG. 14

1500

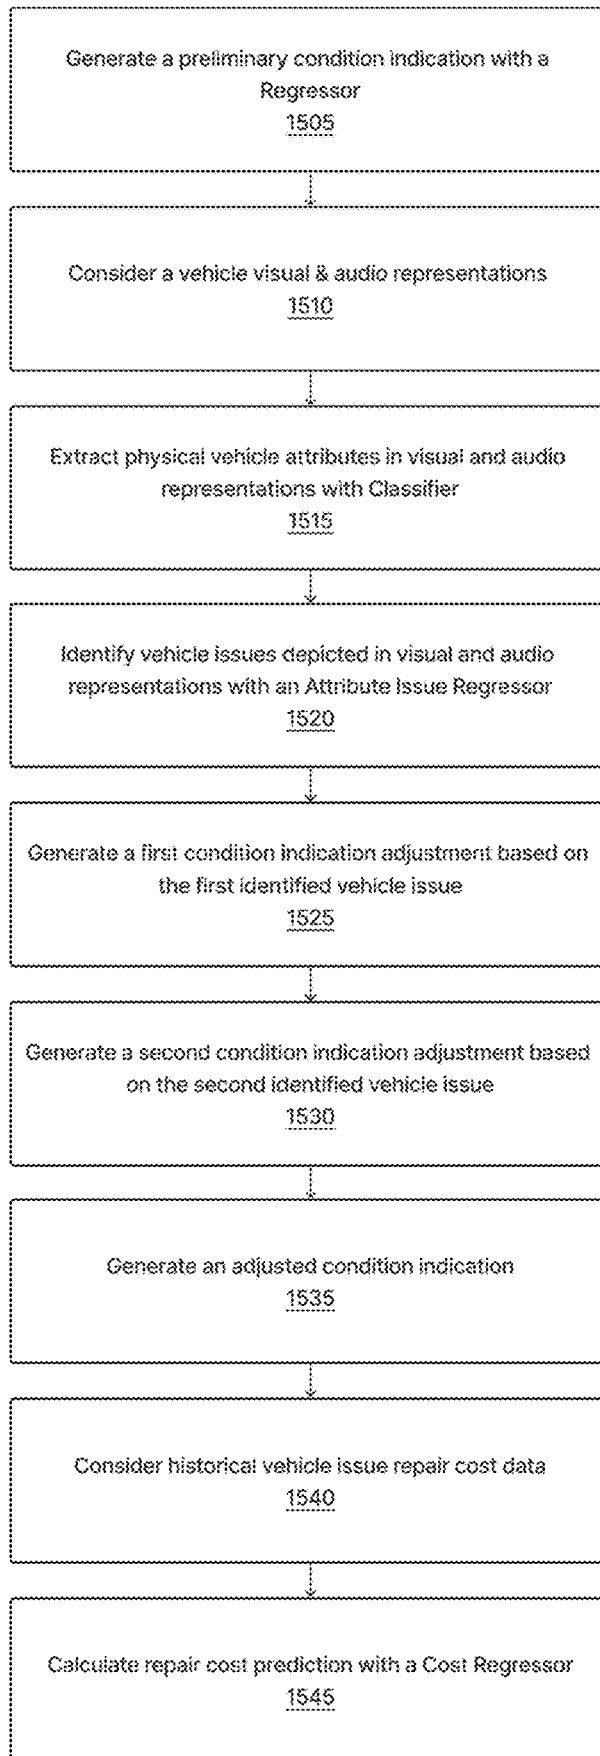
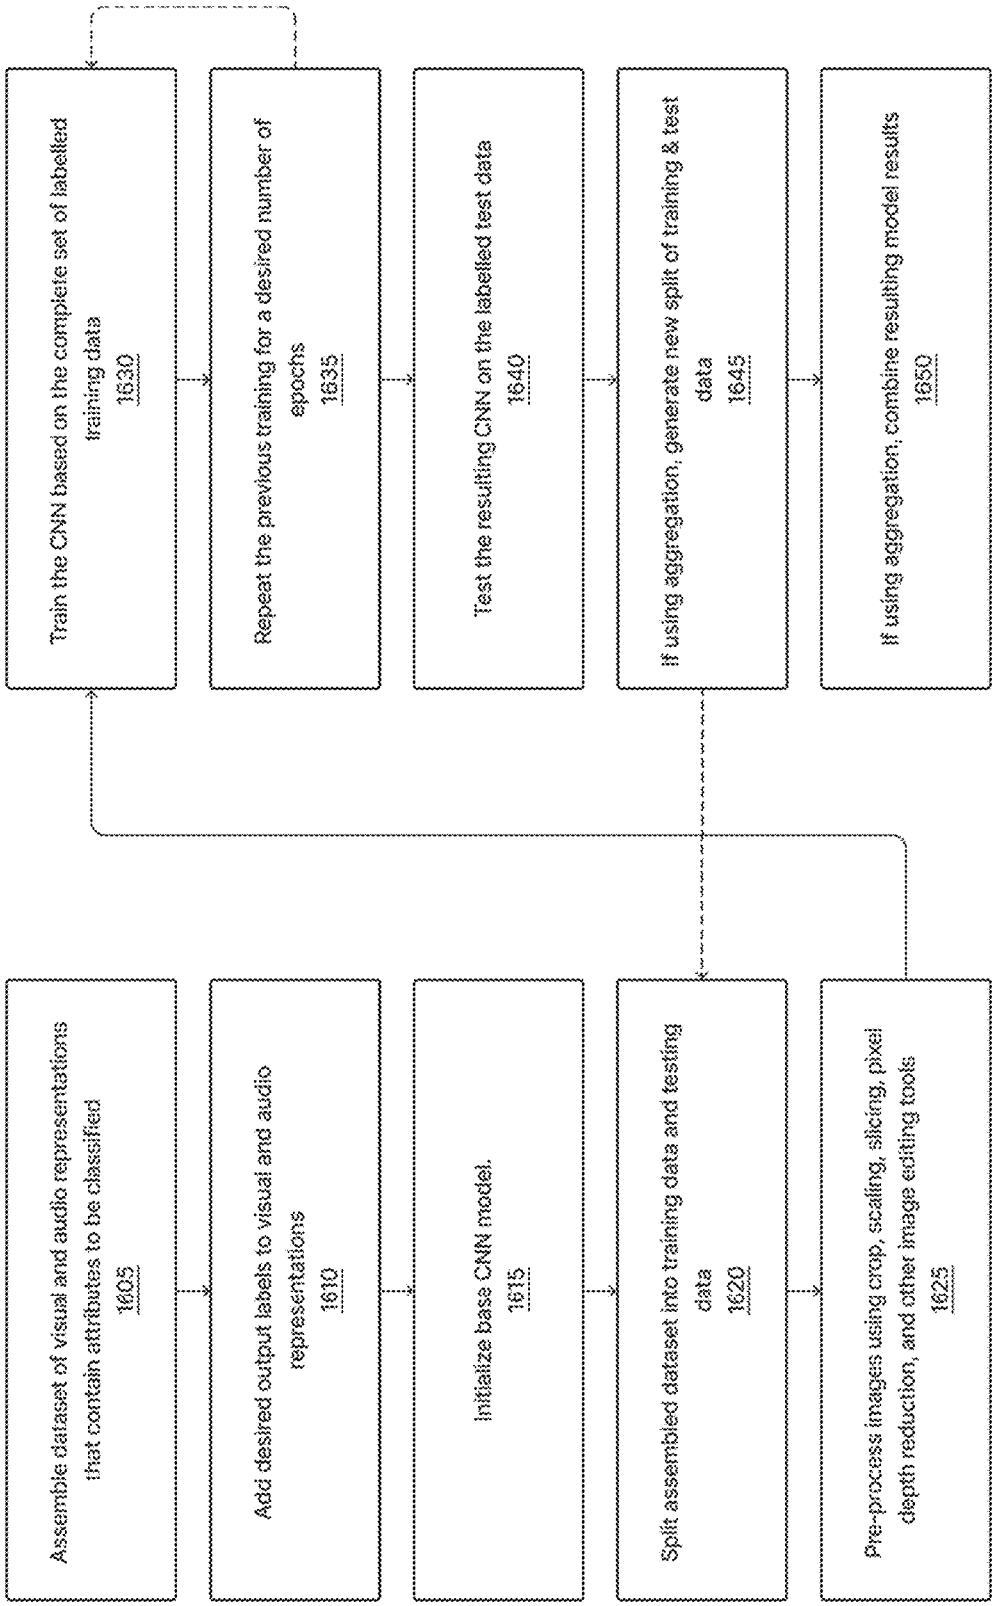


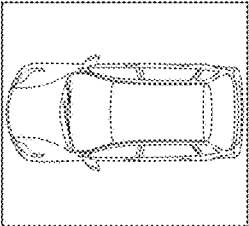
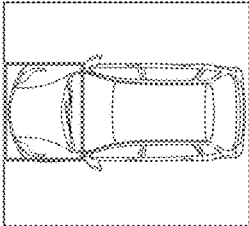
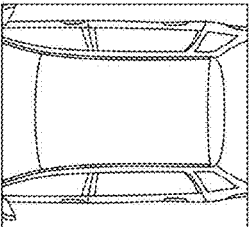
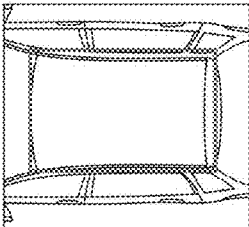
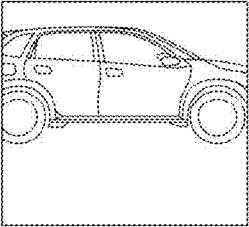
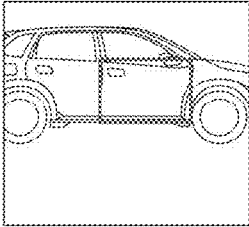
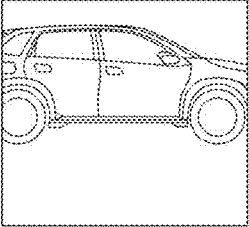
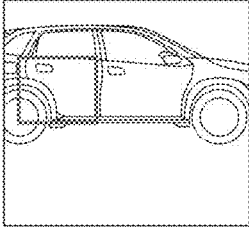
FIG. 15



1600

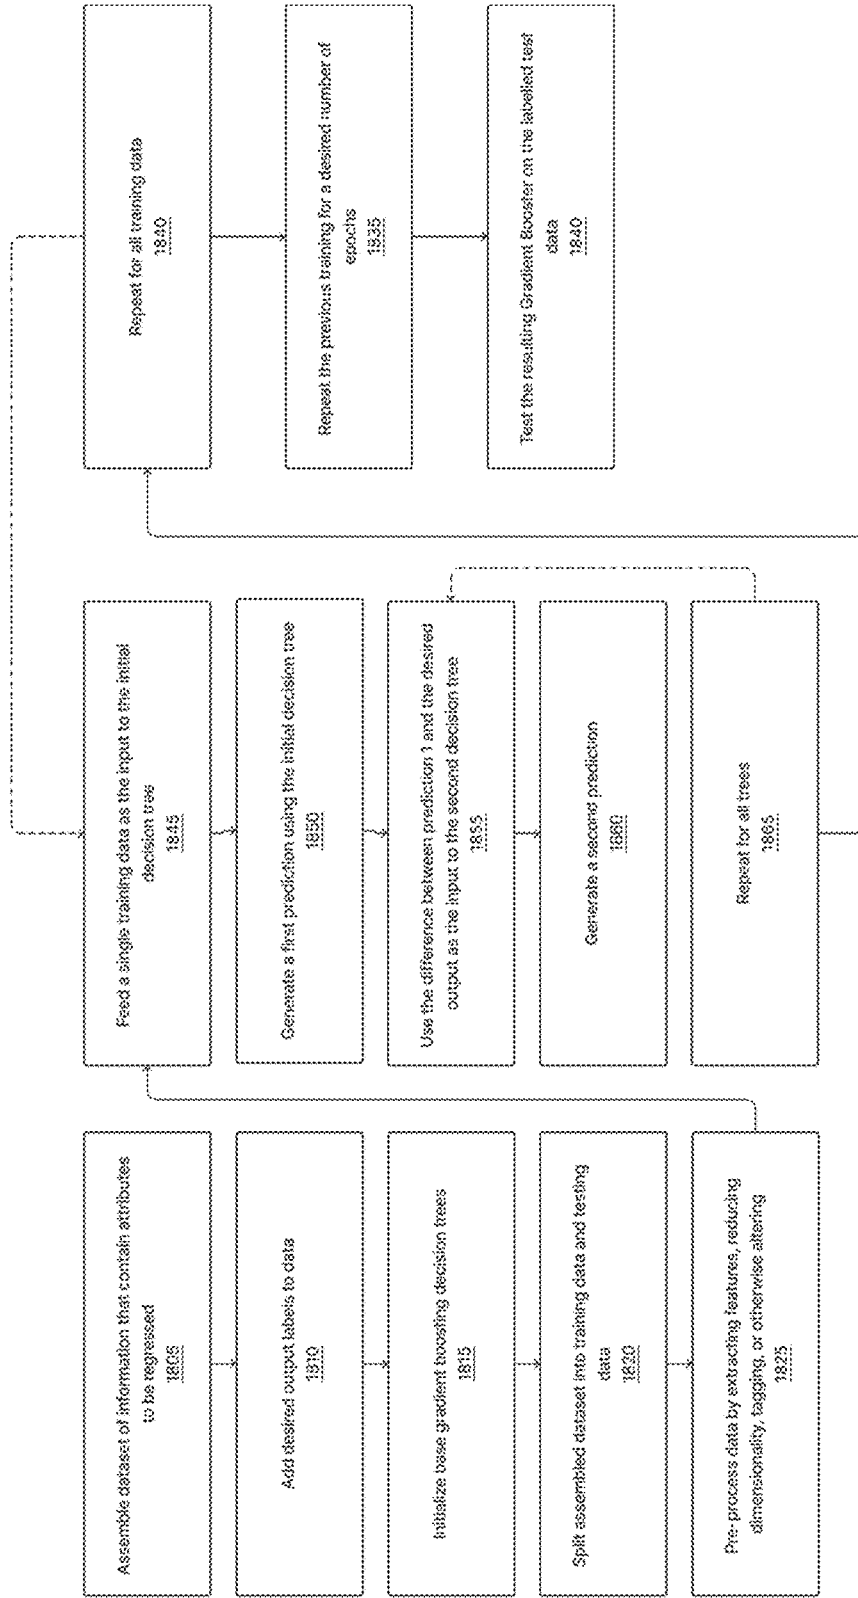
FIG. 16

Example training data for CNN Classifier

	Raw Image	Location	Classification
1.			Hood
2.			Roof
3.			Front right door
4.			Rear right door

1700

FIG. 17



1800

FIG. 18


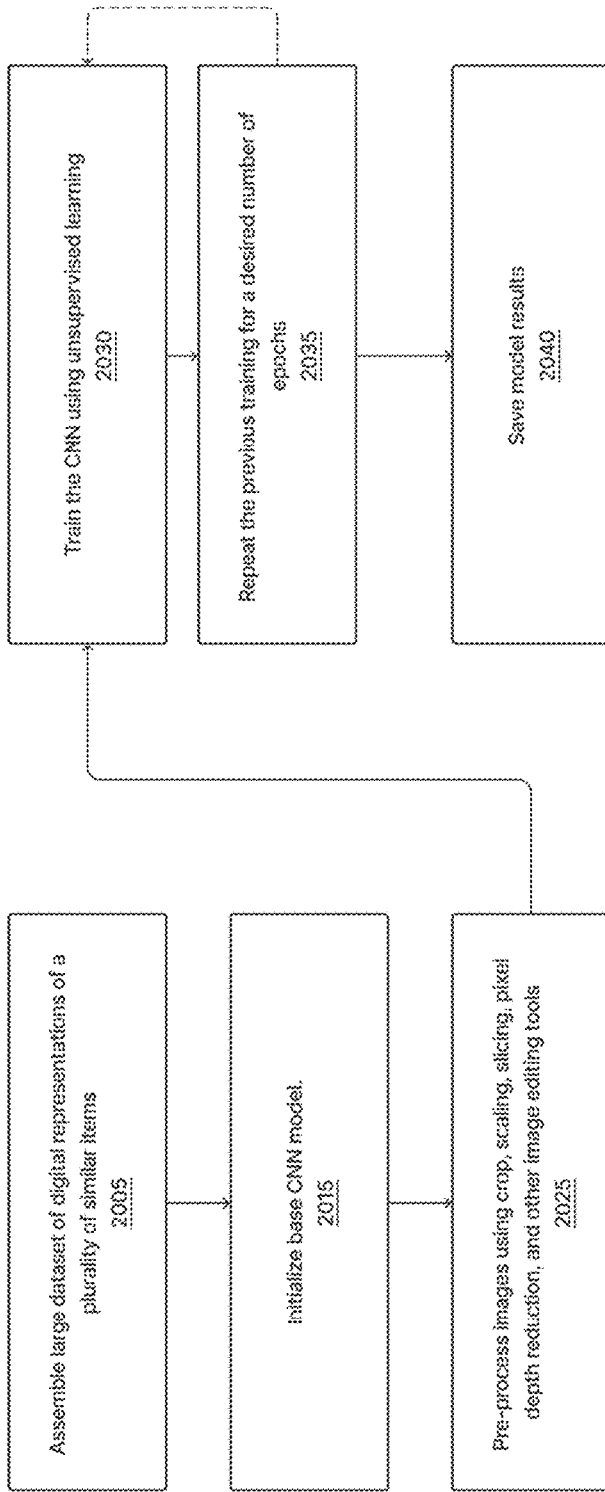


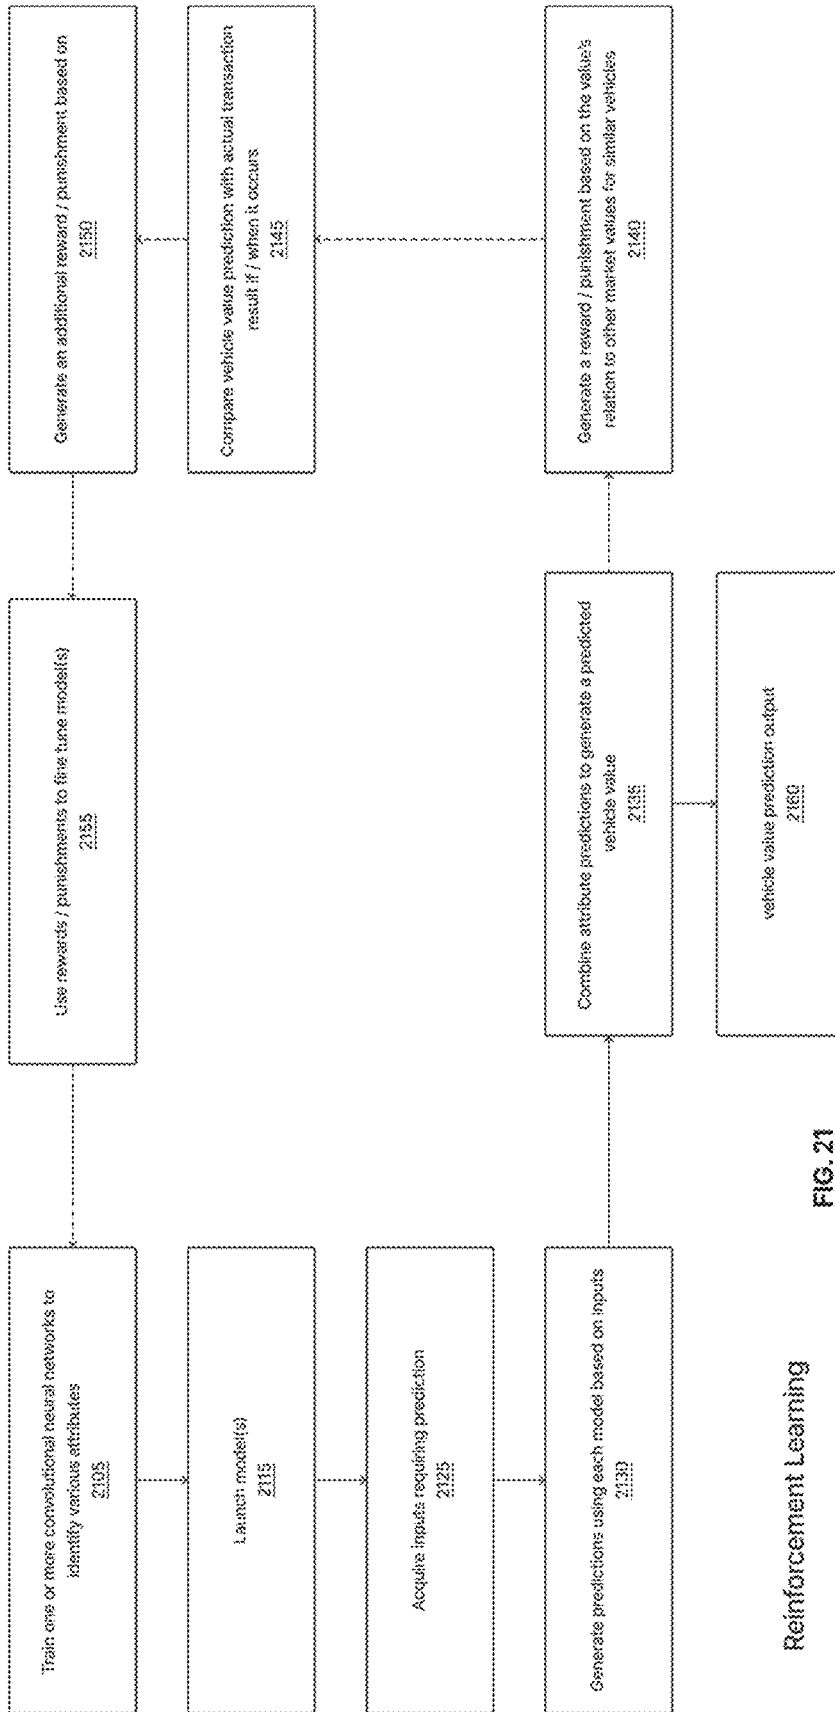
FIG. 19



Unsupervised learning

2000

FIG. 20



Reinforcement Learning

2100

FIG. 21

SYSTEMS AND METHODS FOR VALUING AN ITEM

CROSS-REFERENCE TO RELATED APPLICATIONS

[0001] This application claims priority from U.S. Provisional application No. 63/317,945, filed Mar. 8, 2022, the entire contents of which are hereby incorporated by reference.

BACKGROUND OF INVENTION

[0002] Used items such as vehicles are bought and sold in large volumes on a daily basis. These transactions occur through one or more of a variety of parties including private sellers, auction houses, dealers (also known as resellers), online marketplaces, and private buyers. Traditionally, pricing of used items such as vehicles is determined based on some combination of personal experience, physical inspection, and various pricing guides. The resulting process for a potential buyer/reseller to successfully determine an exact dollar amount for any given vehicle remains manual and unrepeatable. Worse, an organization focused on the buying and selling of vehicles may end up with wildly different expectations of a vehicle's value depending on which of their agents performed the valuation process.

[0003] Currently available service providers offer piecewise solutions to the core problem these potential vehicle buyers have. These providers offer sets of data which may include the following:

[0004] a. Estimated open-market price ranges for a specific vehicle model with emphasis on brackets for abstract vehicle condition such as 'Excellent', 'Average', and 'Rough'.

[0005] b. Physical vehicle evaluation and generation of static 'Condition Reports' which often include a series of questions regarding the condition of vehicle attributes like the presence of dents or scratches.

[0006] The previously noted data is currently available from a variety of service providers and is the foundation for the typical process of valuing a vehicle as it occurs today across thousands of new and used vehicle resellers. While these offerings may solve a piece of the problem for a potential buyer, they leave the real analysis and pricing generation up to the agent performing the assessment. The agents are currently required to have deep automotive background to make informed assessments of the impact of specific physical features on the ultimate value of the vehicle in question. As previously mentioned, this requires organizations employing these agents to maintain rigid, thorough, and expensive training regimens for their employees to attempt to instill consistency in their inventory costs. Yet, due to the abstract—and often un-trackable—decision making occurring at the agent-level, these organizations have little to no insight into inventory pricing variations and how they are effecting their overall business.

SUMMARY OF INVENTION

[0007] In accordance with some embodiments of the present disclosure, a method of valuing an item, comprises: obtaining identification information for the item; obtaining item specifications for the item using the identification information; determining, using a computer processor, a preliminary condition indication based on the item specifications;

obtaining photographs of the item; determining, from the visual representation using a machine learning (ML) engine, a first item attribute; determining, using the computer processor, from the first item attribute, a first condition indication adjustment; determining, using the computer processor, an adjusted condition indication based on the preliminary condition indication and the first condition indication adjustment; obtaining market valuation data of the item based on the identification information; and determining, using the computer processor, a value of the item based on the market valuation data and the adjusted condition indication.

[0008] In accordance with some embodiments of the present disclosure, a non-transitory computer readable medium stores instructions executable by a computer processor for valuing an item, the instructions comprising functionality for: obtaining identification information for the item; obtaining item specifications for the item using the identification information; determining, using a computer processor, a preliminary condition indication based on the item specifications; obtaining a visual representation of the item; determining, from the visual representation using a machine learning (ML) engine, a first item attribute; determining, using the computer processor, from the first item attribute, a first condition indication adjustment; determining, using the computer processor, an adjusted condition indication based on the preliminary condition indication and the first condition indication adjustment; obtaining market valuation data of the item based on the identification information; and determining, using the computer processor, a value of the item based on the market valuation data and the adjusted condition indication.

[0009] In accordance with some embodiments of the present disclosure, a system for valuing an item comprises: a client web interface configured to input identification information for the item and a visual representation of the item; and a computer system configured to: obtain item specifications for the item using the identification information; determine a preliminary condition indication based on the item specifications; determine, from the visual representation using a machine learning (ML) engine, a first item attribute; determine, from the first item attribute, a first condition indication adjustment; determine an adjusted condition indication based on the preliminary condition indication and the first condition indication adjustment; obtain market valuation data of the item based on the identification information; and determine a value of the item based on the market valuation data and the adjusted condition indication.

[0010] In accordance with some embodiments of the present disclosure, a computer-implemented method of training a neural network for item valuation may comprise: collecting a set of digital representations of a plurality of similar items; applying one or more transformations to create a modified set of digital representations of the similar items; creating an unsupervised learning first training set comprising the modified set of digital representations; training the neural network using the unsupervised learning first training set; creating a labeled set of digital representations by labeling a subset of the modified set of digital representations and the collected set of digital representations to identify at least a first attribute and a second attribute; creating a supervised learning second training set comprising the labeled set of digital representations; training the neural network using the supervised learning second training

set; inputting a new set of digital representations to the trained neural network; applying outputs of the trained neural network to a valuation predictor and obtaining valuation results; collecting a set of rewards and/or punishments from the valuation results; and further training the neural network, through reinforcement learning, using the set of rewards and/or punishments.

[0011] Other aspects and advantages of the invention will be apparent from the following description and the appended claims.

BRIEF DESCRIPTION OF DRAWINGS

[0012] FIGS. 1A-1B show non-limiting examples of systems and methods to value an item.

[0013] FIG. 2 shows non-limiting example of an open market price predictor

[0014] FIG. 3 shows non-limiting example of a repair cost predictor

[0015] FIGS. 4-12 show non-limiting examples of systems and methods to value an item, in accordance with one or more embodiments of the present disclosure.

[0016] FIG. 13 shows a computer system in accordance with one or more embodiments of the present disclosure.

[0017] FIGS. 14-16 show non-limiting examples of methods to value an item, in accordance with one or more embodiments of the present disclosure.

[0018] FIG. 17 shows non-limiting example of training data for one or more embodiments of CNN classifiers.

[0019] FIG. 18 shows non-limiting example of a process to train a gradient boosting regressor machine learning algorithm using training data.

[0020] FIG. 19 shows non-limiting example of a client web interface used to generate and acquire user-submitted vehicle details.

[0021] FIG. 20 shows non-limiting example of a process to train a base convolutional neural network using unlabeled training data (unsupervised learning).

[0022] FIG. 21 shows non-limiting example of a process to further train a trained convolutional neural network using feedback including rewards and/or punishments (reinforcement learning).

DETAILED DESCRIPTION

[0023] In the following detailed description of embodiments of the disclosure, numerous specific details are set forth in order to provide a more thorough understanding of the disclosure. However, it will be apparent to one of ordinary skill in the art that the disclosure may be practiced without these specific details. In other instances, well-known features have not been described in detail to avoid unnecessarily complicating the description.

[0024] Throughout the application, ordinal numbers (e.g., first, second, third, etc.) may be used as an adjective for an element (i.e., any noun in the application). The use of ordinal numbers is not to imply or create any particular ordering of the elements nor to limit any element to being only a single element unless expressly disclosed, such as using the terms “before”, “after”, “single”, and other such terminology. Rather, the use of ordinal numbers is to distinguish between the elements. By way of an example, a first element is distinct from a second element, and the first element may encompass more than one element and succeed (or precede) the second element in an ordering of elements.

[0025] In the following description of FIGS. 1-21 any component described with regard to a figure, in various embodiments disclosed herein, may be equivalent to one or more like-named components described with regard to any other figure. For brevity, descriptions of these components will not be repeated with regard to each figure. Thus, each and every embodiment of the components of each figure is incorporated by reference and assumed to be optionally present within every other figure having one or more like-named components. Additionally, in accordance with various embodiments disclosed herein, any description of the components of a figure is to be interpreted as an optional embodiment which may be implemented in addition to, in conjunction with, or in place of the embodiments described with regard to a corresponding like-named component in any other figure.

[0026] It is to be understood that the singular forms “a,” “an,” and “the” include plural referents unless the context clearly dictates otherwise. Thus, for example, reference to “a model parameter” includes reference to one or more of such parameters.

[0027] Terms such as “approximately,” “substantially,” etc., mean that the recited characteristic, parameter, or value need not be achieved exactly, but that deviations or variations, including for example, tolerances, measurement error, measurement accuracy limitations and other factors known to those of skill in the art, may occur in amounts that do not preclude the effect the characteristic was intended to provide.

[0028] It is to be understood that one or more of the steps shown in the flowcharts may be omitted, repeated, and/or performed in a different order than the order shown. Accordingly, the scope disclosed herein should not be considered limited to the specific arrangement of steps shown in the flowcharts.

[0029] In accordance with some embodiments, the systems and methods disclosed herein relate to a technical solution for valuing an item, particularly a used item. In many of the examples in the following description the item valued is a used vehicle (automobile, truck, bus, van, motorcycle, tractor, etc., without limitation) however the scope of the invention is not limited to a particular type of item. The item could, in some embodiments, be a house, antique furniture, a musical instrument, artwork, tools, toys, or any other type of item.

[0030] Systems and Methods to Value an Item

[0031] Valuing an item is often difficult because of how many factors are involved in determining a fair purchase price. Further, when the purchasing party represents a reseller, there is the added complexity of ensuring the item is refurbished to saleable condition. Traditionally, pricing of items is determined based on some combination of personal experience, physical inspection, and various pricing guides. However, such traditional approaches generally require personal expertise and judgement and may not yield consistent results from one valuation to another. Therefore, there is a need for a technical solution that inputs appropriate data, evaluates and validates that data according to established criteria, and provides appropriate valuation guidance.

[0032] Continuing, without limitation, with the example of the item being a used vehicle, as a vehicle ages, its mileage increases, parts need to be replaced, it incurs damage and debris from the road, and more. All of these factors impact its value on the open market and therefore to

a potential buyer or reseller. Major repairs after an accident or other historical events also heavily impact the price a new buyer will pay. Additional effort from paying off a loan, buying out a lease, and managing title transfer means that the financial status of that vehicle also plays a part.

[0033] In accordance with some embodiments, the present disclosure includes a fully-adjusted used vehicle valuation resulting from a system that integrates some or all of these different data and generates a human-readable vehicle condition indication that, along with the original data and its post-processed analysis, can then be applied to the current open market prices of that vehicle.

[0034] In accordance with some embodiments and referring to FIG. 1A, determining the value of an item (110) relies on determining the individual factors involved in its calculation. The formula for determining the exact value of an item (115), otherwise defined as the price a potential buyer is willing to pay, requires knowing the exact value of that item on the open market in saleable condition (100), the cost of repairing the item to saleable condition (105), and the desired profit or gross margin (135) in the case of a reseller. In cases of a non-reseller, the desired profit can be understood to be 0. As there are no universally absolute definitions of value, one must generate predictions based on available data. The method of valuing an item (120) described in FIG. 1B is a real-world representation of the theoretical formula described above (110). An open market price predictor (200) is used to generate a predicted value of the vehicle on the open market (125). Concurrently, or consequently, a repair cost predictor (300) is used to generate a predicted repair cost (130). In cases where the potential buyer aims to resell the vehicle, a desired profit/margin (135) can be included. Generation of the predicted value of a vehicle (140) results from subtracting the predicted repair cost (13) and the desired profit/margin (135) from the predicted value of the vehicle on the open market (125).

[0035] FIG. 2 illustrates an example implementation of an open market price predictor (200), according to some embodiments. To generate an accurate prediction of an item's value on the open market, data sources (205) that contain price-affecting attributes of the item are included as inputs to the open market price regressor (225). In the case of our vehicle value estimation, data sources (205) may include some or all vehicle manufacturers data (210), vehicle history data (215), and historical vehicle market pricing data (220). Vehicle manufacturers' data (210) may include, but is not limited to, vehicle year, make, model, options, specifications, recalls, service statistics, etc. Vehicle history data (215) may include, but is not limited to, ownership history, owner count, vehicle location, crash/accident history, repair history, title information, servicing data, mileage data, etc. Historical vehicle market pricing data (220) may include, but is not limited to, local and national wholesale auction vehicle pricing data, local and national open market (retail) vehicle pricing data, and local and national private party (consumer to consumer) vehicle pricing data. To determine the predicted value of the vehicle on the open market (125), the open market price regressor (225) applies the input data against its trained machine learning regression model to predict an expected price. An example embodiment of a value regressor is illustrated in FIG. 8.

[0036] FIG. 3 illustrates an example implementation of a repair cost predictor (300), according to some embodiments. Generating the predicted repair cost (130) leverages the

combination of determining the issues that need to be fixed, and the costs associated with that. This predictor first determines the variations from saleable condition using the vehicle issue identifier (400), then uses that information to more accurately predict costs using the issue repair cost predictor (500). Additional details on the data provided by the vehicle issue identifier (400) is illustrated in FIG. 4. This embodiment illustrates a prediction method that has visible layers, e.g., sequential steps of generating pieces of the equation. Specifically, the vehicle issue identifier (400) may generate an adjusted condition indication (460) which can then be used by the issue repair cost predictor (500), each utilizing their own implementations of machine learning models or combination thereof. In some embodiments, the adjusted condition indication (460) may contain a human-readable condition score or metadata which may be in the form of a range, quality tier, probability, or other metric. Another embodiment may include a system that takes the input data as described in FIG. 4 and directly generates a predicted repair cost (130) without the visible interim step. In this case, the extraction of the preliminary vehicle condition indication (445), the first condition indication adjustment (450), the second condition indication adjustment (455), and the adjusted condition indication (460) may be represented as layers in a single, or more higher dimensional (depth and or width) machine learning models including each of the corresponding components. In some embodiments this may include a pre-processing layer which acts as the generation of the preliminary vehicle condition indication (445) and then subsequent layers which expand and adjust the condition indication according to the inclusion of digital representations.

[0037] FIG. 4 illustrates an example implementation of a vehicle issue identifier (400), according to some embodiments. Similar to when predicting the open market price, identifying issues with a vehicle may rely on some or all of the following data sources: vehicle manufacturers data (210), user-submitted vehicle details (410), vehicle history (215), service and repair records (420), vehicle images (425), vehicle videos (430), and vehicle audio recordings (435). User-submitted vehicle details (410) may include, but are not limited to, current mileage, physical location, subjective condition evaluation, insights to historical usage, etc. Vehicle images (425) may include, but are not limited to, attribute-labeled and unlabeled images of the exterior, interior, mechanical, and undercarriage from a variety of angles. Vehicle videos (430) may include, but are not limited to, videos depicting the exterior, interior, engine compartment, other mechanical features, and undercarriage of the vehicle while off, idling, and/or in motion. Vehicle audio recordings (435) may include, but are not limited to, audio recordings of the vehicle engine, exhaust, suspension, transmission, and brakes while idling, starting, revving, accelerating, stopping, driving, etc. Some or all of these data are used by the vehicle issue extractor (700) to generate the identified vehicle issues output data (440). More details regarding the generation the identified vehicle issues output data (440) and operation of the vehicle issue extractor (700) are illustrated in FIG. 7. Some or all of the vehicle issue extractor (700) input data are used by the vehicle condition regressor (800) to generate a preliminary vehicle condition indication (445). The identified vehicle issues output data (440) may then be split into corresponding condition indication adjustments, such as a first condition indication adjustment (450) and a second

condition indication adjustment (445). The preliminary vehicle condition indication may be combined with the first and second adjustments to create the adjusted condition indication (460). In some embodiments, the combination may involve adding the respective issues surfaced by the vehicle condition regressor (800) and the vehicle issues extractor (700), in others the combination may involve applying a distribution or other mapping function across one or more of the component condition indications. In some embodiments, the preliminary condition indication (445) and the corresponding adjustments may be structured as dictionaries of vehicle attributes and the corresponding issues such as those illustrated in identified vehicle issues output data in (710). The preliminary vehicle condition indication (445) may be a human-readable representation of current condition of the vehicle with respect to a maximum value of identical to original manufacturers specifications. In some embodiments the vehicle condition indication (445) may be represented as a 0-100 number where 100 is perfect condition.

[0038] FIG. 5 illustrates an example implementation of an issue repair cost predictor (500), according to some embodiments. There are a multitude of strategies for estimating the cost of repairing an item to a certain specification. One approach may include using generic averages for repairing that type of item. For example, the US average vehicle reseller refurbishment cost is currently around \$550 per vehicle. As such, one embodiment of the issue repair cost predictor (500) might use the above broad repair cost average (525) as the sole input to the repair cost regressor (530). Pricing information for repair costs can be as detailed as a line-item list of parts and their parts, along with required labor and its costs, or as broad or more so than the country-wide example provided above. Some, all, or a blended mix of these pricing resolution levels may be included in the repair cost regressor's (530) inputs. FIG. 5 illustrates some embodiments of these variations: the average vehicle repair cost for a specific buyer/reseller (505), a per-vehicle type buyer/reseller average repair cost (510), a per-vehicle model buyer/reseller average repair cost (515), the repair costs of parts & labor for specific repairs, broadly defined or broken down by make and model (520), or broad repair cost averages (525) such as, but not limited to, the average reseller refurbishment cost across the US. The repair cost regressor may leverage similar machine learning models and training as applied in the open market price regressor (225) and vehicle condition regressor (800). Details on the architecture, training, and deployment of such models, in some embodiments, is illustrated as an example in FIG. 8.

[0039] FIG. 6 illustrates an example implementation of a predictor operation mode (60), according to some embodiments. Machine learning-based predictors are often involve multiple steps to deploy between training, testing, and operating. In some embodiments, this may include, but is not limited to, a system including a model image loader (64), a preprocessor & model trainer (61), and a model server (62). The image loader can provide backbone storage for the trained models required to generate predictions. The preprocessor & model trainer (61) can execute the steps necessary to either train a base CNN using labelled data or fine tune a currently trained one as well as clean input data to be predicted on. The model server (62) can perform calculations needed by the trained model such as evaluating convolutional steps in a CNN to generate an output in response

to a prediction request (63). In some embodiments this system may run on a single local server, and in others multiple servers located on local or cloud hosting services like AWS and google cloud.

[0040] FIG. 7 illustrates an example implementation of a vehicle issue extractor (700), according to some embodiments. The vehicle issue extractor (700) leverages input data to generate the identified vehicle issues output data (440). The condition, or presence of issues, of the vehicle can be determined using a multitude of media types, referred to herein generically as a "digital representation". In some embodiments these media may include some or all of vehicle images (425), vehicle videos (430), and vehicle audio recordings (435). Vehicle images (425) are input to a vehicle attribute classifier (900), detailed further in FIG. 9, which leverages a trained convolutional neural network (referred to below as CNN) to generate vehicle images labelled by physical attribute (705). Similarly, vehicle videos (430) and vehicle audio recordings (435) may be processed using attribute classifiers (720 and 725, respectively) to identify and labelled the data by physical attribute. Some or all of these labelled images, digital representations, and service & repair records (420) may be used as input data for the feature attribute issue regressor (1000) to generate the identified vehicle issues output data (440).

[0041] FIG. 8 illustrates an example implementation of a vehicle condition regressor (800), according to some embodiments. The regressor may be a gradient boosting machine learning regression engine which can use successive predictions and their residuals as metadata to generate a forest of prediction trees that can be combined through an ensemble predictor (860). Models such as these can be very effective in extracting non-linear correlations between multiple data sources to generate a prediction output (865). In some embodiments, the data being used could include, but is not limited to, vehicle manufacturers data (210), user-submitted vehicle details (410), vehicle history (215), service and repair records (420), etc. Input data to the vehicle condition regressor (800) may be combined, flattened, filtered, dimensionality reduced, or otherwise manipulated before being used as input data 1 (805) for the regressor to calculate a prediction output (865). A trained gradient boosting model may consist of multiple prediction trees as shown in FIG. 8. Each tree may provide an independent prediction (815, 835, 855) that results from input data (805, 825, 845, respectively). These predictions may combine in an ensemble predictor (860) which then generates a single prediction output (865). In some embodiments an ensemble predictor (860) may use a weighted average result, and in others use more elaborate output layers involving additional regression steps. Tree 2 (830) may be trained using the prediction residuals, difference between predicted value and observed value, of prediction 1 (815), which may then use its prediction residuals based on prediction 2 (835) to train a tree 3 (850). This process may be combined into an N number of prediction trees, where N can be chosen as a standard hyperparameter in publicly available models. In some embodiments, N may be selected by repeatedly training the vehicle condition regressor (800) with varying values of N and then selecting the best performing model, models, or subset of models. An example process for training a regressor of this type in some embodiments is illustrated in FIG. 18.

[0042] FIG. 9 illustrates an example implementation of a vehicle attribute classifier (900), according to some embodiments. As costs to repair an item often rely on the attribute of that item that needs fixing, for instance replacing a bumper is less expensive than replacing an engine, this vehicle attribute classifier (900) generates output data that allows for identified issues to also be identified by physical attribute of the vehicle. This example illustrates the extraction of physical attributes from vehicle images (425) but a similar approach can be applied to other types of digital representations including, for example, vehicle videos (430) and vehicle audio recordings (435). In this embodiment, the vehicle images (425) are input with labels indicating the vehicle section such as interior, exterior and mechanical based on the acquisition method of the vehicles, but other embodiments may apply an additional classification step to go from unlabeled images to section-labelled data or just go directly to attribute labelling by skipping the pre-labelling step. The images of exterior (905) are inputted to the exterior classifier (1100), defined in more detailed in FIG. 11, that leverages a CNN trained to classify the images and whether they feature attributes known to be on the exterior of the vehicle, defined as exterior physical attributes 1 (1006) through attribute N, such as, but not limited to, the hood, bumper, front left door, front right door, etc. The images of interior (910) are inputted to the interior classifier (925), similar in architecture to the exterior classifier (1100) and omitted for brevity, leverages a CNN trained to classify the images and whether they feature attributes known to be on the interior of the vehicle, defined as interior physical attributes 1 (1011) through attribute J, such as, but not limited to, the driver seat, dashboard, headliner, passenger seat, rear seats, trunk, etc. The images of mechanical parts (930) are inputted to the mechanical classifier (930), similar in architecture to the exterior classifier (1100) and omitted for brevity, that leverages a CNN trained to classify the images and whether they feature attributes known to illustrate mechanical components of the vehicle, defined as mechanical attributes 1 (1016) through attribute G, such as, but not limited to, the engine, brakes, exhaust, suspension, etc. The results of the exterior classifier (1100), interior classifier (925), and mechanical classifier (930) combine to create the vehicle images labelled by physical attribute (705). This data contains all or a subset of the vehicle images (425), which may or may not have been edited, cropped, sliced, resized or had any other alteration, each with attribute labels attached to understand what section of the vehicle an individual image represents.

[0043] FIG. 10 illustrates an example implementation of a feature attribute issue regressor (1000), according to some embodiments. This regressor may combine raw or pre-processed vehicle specification, history, image, video, and audio data to predict and extract physical vehicle issues associated with the vehicle being valued. These predictions, which may be generated by some or all of the aggregate exterior regressor (1005), the aggregate interior regressor (1010), and the aggregate mechanical regressor (1015) may be combined into an identified vehicle issues output data (440) set describing the identified issues and the corresponding physical attribute affected. An aggregate exterior regressor (1005) may use some or all of the provided service & repair records (420), vehicle images labelled by physical attribute (705), vehicle videos labelled by physical attribute (710), vehicle audio recordings labelled by physical attribute

(715) to generate classifications and regression predictions (1008) of exterior issues (1007) contained in those digital representations. These issues are evaluated for each exterior physical attribute (1006). Exterior issues (1007) to extract may include, but are not limited to, the existence of the object (eg. door panel), minor scratches, major scratches, sizes of scratches, paint chipping, rust, cracked glass, non-functional items like lightbulbs, dimness of lights, or any other physical deviation from manufacturing quality. An aggregate interior regressor (1010) may use some or all of the provided service & repair records (420), vehicle images labelled by physical attribute (705), vehicle videos labelled by physical attribute (710), vehicle audio recordings labelled by physical attribute (715) to generate classifications and regression predictions (1008) of interior issues (1012) contained in those digital representations. These issues are evaluated for each interior physical attribute (1011). Interior issues (1012) to extract may include, but are not limited to, the existence of the object (eg. front seat), minor scratches, major scratches, seat tears, faded material, general stains, smoking stains, non-functional items like lightbulbs or other electronics, dimness of lights, or any other physical deviation from manufacturing quality. An aggregate mechanical regressor (1015) may use some or all of the provided service & repair records (420), vehicle images labelled by physical attribute (705), vehicle videos labelled by physical attribute (710), vehicle audio recordings labelled by physical attribute (715) to generate classifications and regression predictions (1008) of mechanical issues (1017) contained in those digital representations. These issues are evaluated for each mechanical attribute (1016). Mechanical issues (1017) to extract may include, but are not limited to, the existence of the object (eg. broken suspension mechanism), road rash on wheels or tires, tread wear of tires, visible wear on brakes, dripping or flowing liquids, variations in engine noise, variations in exhaust noise, or any other physical deviation from manufacturing quality. In some embodiments, a single, broader-trained machine learning algorithm may be trained to act as a combination of the individual aggregate regressor types. In this embodiment, there may be a single aggregate regressor which can predict the presence of issues on the interior, exterior, mechanical portions, or anywhere else on the vehicle.

[0044] FIG. 11 illustrates an implementation of an exterior classifier (1100), according to some embodiments. The exterior classifier processes visual representations and generates predictions of which ones include attributes for which the CNN has been trained. These attributes may include, but are not limited to, a hood, a bumper, a left headlight, a left door panel, a roof, a wheel, etc. The initial convolution layer (1105) transforms the input image layer by convolving it with a kernel, given a defined stride and padding. In this embodiment the kernel is (3x3), the stride 1, padding set as 'same', and the filter depth k, but other embodiments may configure these differently, or combine variations of such to generate multiple model possibilities from which to extract the best. The input image begins with dimensions of (AxAx1) in this embodiment, but may contain additional dimensionalities in other embodiments. The output of the convolution layer (1105) is a representation of the input layer images but with the pixels having been filtered by the kernel image. The depth of the layer has expanded by k as a result of the application of k filters on the input layer. The pooling layer (1110) applies a simple size reduction step to improve

training and prediction speed. In this embodiment, a max pooling stage of (2×2) is applied which reduces the total pixel count by 2×. Additional convolution layers and pooling layers (1111) may or may not be applied to improve overall model accuracy and performance. After the additional convolutional & pooling layers (1111) the image set is flattened into a 1-dimensional object which can then be passed through traditional neural network neurons in the fully connected layers (1120). This embodiment illustrates 2 fully connected internal layers, but other embodiments could include 0, 1, or more of these layers. The output layer (1125) is the sigmoid activation layer which converts our hidden layers into a classification object output. The output classification (1130) is the result of which, if any of the trained attributes for which the CNN is searching, exists in the input image.

[0045] While in this embodiment the convolutional neural network (CNN) with output classification illustrates the application to classifying exterior attributes of a vehicle, the architecture can equally be applied to alternative classifying categories such as, but not limited to, attributes of the interior of a vehicle, attributes of the mechanical sections of a vehicle, uncategorized attributes of a vehicle (e.g., the ability to classify the attribute directly regardless of if the attribute is inside or outside of the vehicle), etc.

[0046] FIG. 12 illustrates an implementation of an exterior regressor (1200), according to some embodiments. This embodiment illustrates the exterior regressor as a CNN with output regression, but other embodiments may employ one of any number of additional machine learning model types that are also suited to generating an output value (vs. a classification). The exterior regressor processes visual representations and generates predictions of the intensity, count, distribution, size, etc. of attributes included in said visual representation. These attributes may include, but are not limited to, number of issues like scratches found, number of dents, size of dents, quality of paint, wear on tires, etc. The initial convolution layer (1205) transforms the input image layer by convolving it with a kernel, given a defined stride and padding. In this embodiment the kernel is (3×3), the stride 1, padding set as 'same', and the filter depth k, but other embodiments may configure these differently, or combine variations of such to generate multiple model possibilities from which to extract the best. The input image begins with dimensions of (A×A×1) in this embodiment, but may contain additional dimensionalities in other embodiments. The output of the convolution layer (1205) is a representation of the input layer images but with the pixels having been filtered by the kernel image. The depth of the layer has expanded by k as a result of the application of k filters on the input layer. The pooling layer (1210) applies a simple size reduction step to improve training and prediction speed. In this embodiment, a max pooling stage of (2×2) is applied which reduces the total pixel count by 2×. Additional convolution layers and pooling layers (1211) may or may not be applied to improve overall model accuracy and performance. After the additional convolutional & pooling layers (1211) the image set is flattened into a 1-dimensional object which can then be passed through traditional neural network neurons in the fully connected layers (1220). This embodiment illustrates 2 fully connected internal layers, but other embodiments could include 0, 1, or more of these layers. The output layer (1225) is the sigmoid activation layer that converts the hidden layers into a classification

object output. The output classification (1230) is the result of which, if any, of the trained attributes for which the CNN is searching, exists in the input image.

[0047] While in this embodiment the convolutional neural network (CNN) with output classification illustrates the application to identify issues of exterior attributes of a vehicle, the architecture can equally be applied to alternative identification categories such as, but not limited to, interior attribute issues, mechanical attribute issues, etc.

[0048] FIG. 13 shows a computer system in accordance with one or more embodiments. FIG. 13 further depicts a block diagram of a computer system (602) used to provide computational functionalities associated with described algorithms, methods, functions, processes, flows, and procedures as described in this disclosure. The illustrated computer (602) is intended to encompass any computing device such as a server, desktop computer, laptop/notebook computer, wireless data port, smart phone, personal data assistant (PDA), tablet computing device, one or more processors within these devices, or any other suitable processing device, including both physical or virtual instances (or both) of the computing device. Additionally, the computer system (602) may include a computer system that includes an input device, such as a keypad, keyboard, touch screen, or other device that can accept user information, and an output device that conveys information associated with the operation of the computer system (602), including digital data, visual, or audio information (or a combination of information), or a graphical user interface (GUI).

[0049] The computer system (602) can serve in a role as a client, network component, a server, a database or other persistency, or any other component (or a combination of roles) of a computer system (602) for performing the subject matter described in the instant disclosure. The illustrated computer system (602) is communicably coupled with a network (630). In some implementations, one or more components of the computer system (602) may be configured to operate within environments, including cloud-computing-based, local, global, or other environment (or a combination of environments).

[0050] At a high level, the computer system (602) is an electronic computing device operable to receive, transmit, process, store, or manage data and information associated with the described subject matter. According to some implementations, the computer system (602) may also include or be communicably coupled with an application server, e-mail server, web server, caching server, streaming data server, business intelligence (BI) server, or other server (or a combination of servers).

[0051] The computer system (602) can receive requests over network (630) from a client application (for example, executing on another computer system (602) and responding to the received requests by processing the said requests in an appropriate software application. In addition, requests may also be sent to the computer system (602) from internal users (for example, from a command console or by other appropriate access method), external or third-parties, other automated applications, as well as any other appropriate entities, individuals, systems, or computers.

[0052] Each of the components of the computer system (602) can communicate using a system bus (603). In some implementations, any or all of the components of the computer system (602), both hardware or software (or a combination of hardware and software), may interface with each

other or the interface (604) (or a combination of both) over the system bus (603) using an application programming interface (API) (612) or a service layer (613) (or a combination of the API (612) and service layer (613)). The API (612) may include specifications for routines, data structures, and object classes. The API (612) may be either computer-language independent or dependent and refer to a complete interface, a single function, or even a set of APIs. The service layer (613) provides software services to the computer system (602) or other components (whether or not illustrated) that are communicably coupled to the computer system (602). The functionality of the computer system (602) may be accessible for all service consumers using this service layer. Software services, such as those provided by the service layer (613), provide reusable, defined business functionalities through a defined interface. For example, the interface may be software written in JAVA, C++, or other suitable language providing data in extensible markup language (XML) format or another suitable format. While illustrated as an integrated component of the computer system (602), alternative implementations may illustrate the API (612) or the service layer (613) as stand-alone components in relation to other components of the computer system (602) or other components (whether or not illustrated) that are communicably coupled to the computer system (602). Moreover, any or all parts of the API (612) or the service layer (613) may be implemented as child or sub-modules of another software module, enterprise application, or hardware module without departing from the scope of this disclosure.

[0053] The computer system (602) includes an interface (604). Although illustrated as a single interface (604) in FIG. 6, two or more interfaces (604) may be used according to particular needs, desires, or particular implementations of the computer system (602). The interface (604) is used by the computer system (602) for communicating with other systems in a distributed environment that are connected to the network (630). Generally, the interface (604) includes logic encoded in software or hardware (or a combination of software and hardware) and operable to communicate with the network (630). More specifically, the interface (604) may include software supporting one or more communication protocols associated with communications such that the network (630) or interface's hardware is operable to communicate physical signals within and outside of the illustrated computer system (602).

[0054] The computer system (602) includes at least one computer processor (605). Although illustrated as a single computer processor (605) in FIG. 6, two or more processors may be used according to particular needs, desires, or particular implementations of the computer system (602). Generally, the computer processor (605) executes instructions and manipulates data to perform the operations of the computer system (602) and any algorithms, methods, functions, processes, flows, and procedures as described in the instant disclosure.

[0055] The computer system (602) also includes a memory (606) that holds data for the computer system (602) or other components (or a combination of both) that can be connected to the network (630). For example, memory (606) can be a database storing data consistent with this disclosure. Although illustrated as a single memory (606) in FIG. 6, two or more memories may be used according to particular needs, desires, or particular implementations of the com-

puter system (602) and the described functionality. While memory (606) is illustrated as an integral component of the computer system (602), in alternative implementations, memory (606) can be external to the computer system (602).

[0056] The application (607) is an algorithmic software engine providing functionality according to particular needs, desires, or particular implementations of the computer system (602), particularly with respect to functionality described in this disclosure. For example, application (607) can serve as one or more components, modules, or applications. Further, although illustrated as a single application (607), the application (607) may be implemented as multiple applications (607) on the computer system (602). In addition, although illustrated as integral to the computer system (602), in alternative implementations, the application (607) can be external to the computer system (602).

[0057] There may be any number of computers (602) associated with, or external to, a computer system (602), wherein each computer (602) communicates over network (630). Further, the term "client," "user," and other appropriate terminology may be used interchangeably as appropriate without departing from the scope of this disclosure. Moreover, this disclosure contemplates that many users may use one computer system (602), or that one user may use multiple computer systems (602).

[0058] FIG. 14 illustrates an example process (1400) for calculating the value of a vehicle, according to some embodiments. In step 1405, process 1400 can collect vehicle condition, history and financial information. In some embodiments, step 1405 may collect information from the vehicle's owner, and in others from external data sources, and still others a combination of the two. In step 1415, process 1400 may consider additional vehicle information including but not limited to some or all of vehicle manufacturers data (210), vehicle history data (215), and historical vehicle market pricing data (220). Vehicle manufacturers' data (210) may include, but is not limited to, vehicle year, make, model, options, specifications, recalls, service statistics, etc. Vehicle history data (215) may include, but is not limited to, ownership history, owner count, vehicle location, crash/accident history, repair history, title information, servicing data, mileage data, etc. Historical vehicle market pricing data (220) may include, but is not limited to, local and national wholesale auction vehicle pricing data, local and national open market (retail) vehicle pricing data, and local and national private party (consumer to consumer) vehicle pricing data. Process 1400 may use these data in step 1420 and process 1500 to generate predictions for the open market price with a Regressor and the vehicle repair cost, respectively. Step 1420 may use the open market price regressor (225), a trained multi-layer CNN with an output regression layer, to evaluate the inputted vehicle information and generate a predicted price. Process 1500, described in detail in FIG. 15, may use the inputted vehicle information to predict the vehicle repair cost. Step 1425 may consider profit/gross margin preferences of the prospective buyer to adjust the perceived vehicle value. This value may contain an exact value, a sliding scale, a classification engine, or other implementations. Cases where the prospective buyer is not a reseller, the desired profit margin may be considered to be 0. Step 1430 may calculate the resulting value of the vehicle by combining the open market price, predicted vehicle repair cost, and the profit/gross margin input.

[0059] FIG. 15 illustrates an example process (1500) for predicting the repair cost for a vehicle, according to some embodiments. In step 1505, process 1500 may use a condition regressor to generate a preliminary condition indication using the vehicle information acquired in process 1400. In step 1510, process 1500 can collect a digital representation or representations of the vehicle. The digital representation may be a visual representation that may include, but is not limited to, a photograph, a set of photographs, a video, or set of videos depicting various angles of the vehicle. In some embodiments, the digital representation may include an audio representation or representations. An audio representation may include, but is not limited to, audio recordings of the vehicle engine, exhaust, suspension, transmission, and brakes while idling, starting, revving, accelerating, stopping, driving, etc. Step 1515 may use a CNN-based attribute classifier to extract physical vehicle attributes from the visual representation or representations or audio representation or representations. Physical attributes may include but are not limited to the hood, the left door, the vehicle frame, the dashboard, the entertainment system, the front seat, the headliner, the engine, the transmission, the brakes, the exhaust, etc. Step 1520 may use an attribute issue regressor to identify vehicle issues depicted in some or all of the visual and or audio representations. Step 1525 may generate a first condition indication adjustment based on the first identified vehicle issue in Step 1520. In some embodiments this may include a representation of the physical attribute affected, such as the door panel, and the corresponding issue, such as a major dent. In other embodiments, this may be a discrete value which indicates a calculated score such as 75 out of a potential 100 score. Step 1530 may generate a second condition indication adjustment based on the second identified vehicle issue in Step 1520. Step 1535 may generate an adjusted condition indication based on the combination of the preliminary condition indication, the first, and the second condition indication adjustments. Step 1540 may consider historical vehicle issue repair cost data to include in the repair cost calculations in step 1545. Historical vehicle issue repair costs may include but is not limited to, repair costs from the buyer, repair costs by vehicle make or type, aggregated repair cost averages, etc. Step 1545 may use a cost regressor to predict the total vehicle repair cost using the result of step 1540.

[0060] FIG. 16 illustrates an example process (1600) for training a convolutional neural network, according to some embodiments. Step 1605 may include assembly of a dataset of visual and or audio representations that contain attributes to be classified. In some embodiments a digital representation may include a visual representation that may include but is not limited to an image of a vehicle door, a vehicle roof, a vehicle dashboard, or a video of a vehicle door, a vehicle exterior, a vehicle interior. In some embodiments a digital representation may include an audio representation that may include representations as described in step 1510. Step 1610 may add desired output labels to visual and audio representations assembled in step 1605. These labels may include, but are not limited to, indicating the presence of a physical vehicle attribute such as the front left door or trunk, the presence of an issue or imperfection such as a dent, scratch, odd noise in the exhaust or engine, the intensity of an imperfection such as the size of a dent, the length of a scratch, a rating for cleanliness of the interior, or any other attribute. Step 1615 may include an initialization of a base

convolutional neural network (CNN). In some embodiments, this base CNN may be a pre-trained image model provided by services like Google AutoML or AWS SageMaker. Pre-trained models offer many performance and timing benefits as they have been configured to specialize in image and video attribute classification and regression based on training cycles on billions of images on the internet. Step 1620 may include splitting the assembled dataset from step 1605 and 1610 into separate sections: one for training and the other testing. Training data allows for configuration of the optimizing of the CNN weights and activation energies while testing data provides insight into the likely accuracy of the resulting model. Step 1625 may include preprocessing of images using crop, scaling, slicing, pixel depth reduction, or any other image editing tools and mechanisms. Step 1630 may include training the CNN using the complete set of labelled training data. Step 1635 may include repeating step 1630 for a desired number of training epochs to fully propagate weight configurations through the layers of the model. In some embodiments, the number of epochs can be varied to determine the value which generates them optimum resulting CNN model. Step 1640 may include testing the resulting CNN model or models on the labelled test data. Step 1645 may include generating a new split of training & test data and then performing Steps 1620-1640 again to create a second, third or more trained CNN models. Step 1650 may include combining the resulting plurality of CNN models in the case that aggregation was used in Step 1645. In some embodiments, the combination may include selecting the highest performing model, averaging the results of multiple models, performing addition random forest analysis, or any other type of model combination strategies.

[0061] FIG. 17 illustrates an example table of example training data for a CNN classifier, according to some embodiments. A raw image can be associated with a classification which labels it as a specific attribute of a vehicle, and a location at which the attribute can be found. In some embodiments, this may include the classification of external vehicle body panels and their location on the raw image. In some embodiments, a video may be included in training data which could include similar labels on a frame by frame basis, an overall basis, some fraction thereof, or other configurations. In some embodiments the training data may include audio representations of the vehicle. Labelling for an audio representation may include, but are not limited to, a classification of the presence of an issue (e.g., a spark plug problem or lack of catalytic converters), a representation of the time at or time range of which an issue is present, or other configurations.

[0062] FIG. 18 illustrates an example process (1800) for training a gradient boosting classifier or regressor, according to some embodiments. A gradient boosting classifier or regressor have been illustrated as possible embodiments of the vehicle condition regressor (800) and the repair cost regressor (530). Step 1805 may involve assembling a dataset that contains attributes to be regressed, or in some embodiments, classified. Step 1810 may involves adding desired output labels to data. Step 1815 may involving initializing base gradient boosting decision trees. Step 1820 may split an assembled dataset into training data and testing data. Step 1825 can pre-process assembled data by extracting features, reducing dimensionality, tagging, or otherwise altering it. Step 1845 can use an element of the data output by step 1825 as a training data input to the initial decision tree. Step 1850

can generate a first prediction using the initial decision tree. Step **1855** can use the difference between prediction **1** and the desired output as the input to the second decision tree. Step **1860** can generate a second prediction using the residual of step **1855**. Step **1865** can repeat this propagation sequence of residuals through all defined trees in the model. Step **1840** can repeat the full tree training process for the rest of the training data set. Step **1835** can repeat the training process across all training data for a given number of times, otherwise known as epochs. Step **1840** can test the resulting gradient boosting algorithm against the labelled test data.

[0063] FIG. **19** shows a client web interface (**1900**) in accordance with one or more embodiments. A client web interface (**1900**) can use a yin number, license plate, or other identification method to determine basic year, make, model information about a user's vehicle through the VIN number or license plate form (**1906**). A vehicle condition report form (**1925**) can include sub forms to gather additional information about the vehicle from a user. A vehicle basics form (**1910**) can allow users to submit mileage, location, color, and/or other basic vehicle information. A vehicle details form (**1915**) can allow users to submit further details about the vehicle such as the lien information, any modifications that have been made, and a vehicle/crash history overview. A vehicle condition form (**1920**) can allow users to submit detailed reporting of the current physical and mechanical condition of the vehicle. This may include, but is not limited to, ranking features on a sliding scale (eg. 1-10), identifying specific damage and/or issue and the vehicle attributes they affect, written explanations of the damage, digital representations, etc. The combined aggregation of the data supplied by and pulled on behalf of the use, may be used to generate the user submitted vehicle details (**410**) object.

[0064] FIG. **20** illustrates an example process (**2000**) for training a convolutional neural network using unsupervised training, according to some embodiments. Step **2005** can involve assembly of a large dataset of digital representations of a plurality of similar items. Step **2015** can initialize a base CNN model. In some embodiments, this base CNN may be a pre-trained image model provided by services like Google AutoML or AWS SageMaker. Pre-trained models offer many performance and timing benefits as they have been configured to specialize in digital representations based on training cycles on billions of images on the internet. Step **2025** may pre-process images using crop, scaling, slicing, pixel depth reduction, and or other image editing tools. Step **2030** can train the CNN using unsupervised learning. Step **2035** can repeat the unsupervised learning step of **2030** for a number of epochs. In some embodiments, epoch count can be left as an industry standard, and in others picked based on a test set of varying epoch counts and performing basic gradient descent. Step **2040** can save the resulting trained model for use in classification, regression, or whatever the model was trained to predict.

[0065] FIG. **21** illustrates an example process (**2100**) for further training a convolutional neural network using reinforcement learning, according to some embodiments. Step **2105** can train one or more convolutional neural networks to identify various attributes. In some embodiments these attributes may be vehicle features displayed in digital representations. In others it may be pricing predictions based on time-wise historical market pricing data. Step **2115** can launch these trained models into an operational mode where it or they are used in their primary purpose which is to

predict outputs based on given inputs. Step **2125** can involve acquiring the inputs requiring prediction by the trained models. Step **2130** can generate predictions from the respective trained models based on each of their respective data inputs. In some embodiments, these inputs may all be shared, in others, some inputs may be used only by a subset of the trained models while other inputs may be used by yet another subset of the trained models. Step **2135** can combine attribute predictions by individual models to generate a predicted vehicle value. Step **2160** can output this vehicle value prediction to other parts of a system. Step **2140** can generate a reward/punishment based on the value's relation to other market values for similar vehicles. In some embodiments this may utilize the mean squared error from average market price for a given year/make/model vehicle to then scale the distribution of the reward/punishment. An example might be a +1 reward for remaining in the first standard deviation, while a -5 could be used when outside the third standard deviation of mean squared error distributions. Step **2145** can compare the vehicle value prediction with actual transaction result if/when it occurs. In some embodiments this may include, but is not limited to, publicly available transaction data, future listing price of the vehicle when resold, or internally determined prices based on any interaction or direct transactions with the owner of the vehicle. Step **2150** can generate an additional reward/punishment based on the variation from observed transaction data. In some embodiments this could assign a +5 for predicted values that fall within 5% of the observed transaction price and a -5 for those that fall outside of 20% range, use a sliding scale gradient, use a discrete bracket system, or any number of other methodologies for assigning reward/punishment values. Step **2155** can use the assigned rewards/punishments to then feed back into the individual CNN's of step **2105**. In some embodiments, this reward mechanism can provide a secondary layer of training information that is separate from the original information the model has access to. For instance, in some embodiments one of the models may be trained to predict the existence of dents from digital representations. A labelled training dataset of where and/or if there are dents in an image has no insight into the ultimate value of the item being detected. The physical system, however, may be correlated in that value of an item can be greatly affected if there are big dents everywhere. As such a reward/punishment mechanism can allow for greater dimensionality of training and ultimately higher accuracy than with a standard supervised training approach.

[0066] While the invention has been described with respect to a limited number of embodiments, those skilled in the art, having benefit of this disclosure, will appreciate that other embodiments can be devised which do not depart from the scope of the invention as disclosed herein. Accordingly, the scope of the invention should be limited only by the attached claims.

What is claimed is:

1. A method of valuing an item, comprising:
 - obtaining identification information for the item;
 - obtaining item specifications for the item using the identification information;
 - determining, using a computer processor, a preliminary condition indication based on the item specifications;
 - obtaining a digital representation of the item;
 - determining, from the digital representation using a machine learning (ML) engine, a first item attribute;

- determining, using the computer processor, from the first item attribute, a first condition indication adjustment; determining, using the computer processor, an adjusted condition indication based on the preliminary condition indication and the first condition indication adjustment; obtaining market valuation data of the item based on the identification information; and determining, using the computer processor, a value of the item based on the market valuation data and the adjusted condition indication.
2. The method of claim 1, further comprising: determining, from the digital representation using the ML engine, a second item attribute; determining, using the computer processor, from the second item attribute, a second condition indication adjustment; and determining, using the computer processor, the adjusted condition indication based on the preliminary condition indication, the first condition indication adjustment, and the second condition indication adjustment.
3. The method of claim 1, wherein the determining, from the digital representation using the ML engine, the first item attribute comprises:
- obtaining a first plurality of digital representations of first similar items;
 - labelling the first plurality of digital representations of first similar items with first training attributes; and
 - training the ML engine with the labelled first plurality of digital representations.
4. The method of claim 3, wherein the determining, from the digital representation using the ML engine, the second item attribute comprises:
- obtaining a second plurality of digital representations of second similar items;
 - labelling the second plurality of digital representations of second similar items with second training attributes; and
 - training the ML engine with the labelled second plurality of digital representations.
5. The method of claim 1, wherein the digital representation comprises a visual representation.
6. The method of claim 1, wherein the digital representation comprises an audio representation.
7. The method of claim 1, further comprising: obtaining history data of the item; wherein determining the preliminary condition indication is further based on the history data.
8. The method of claim 1, further comprising: obtaining ownership information of the item; wherein determining the preliminary condition indication is further based on the ownership information.
9. The method of claim 1, wherein the item is a used vehicle.
10. The method of claim 1, wherein the first item attribute comprises at least one of: existence of dents; existence of scratches; existence of interior stains; condition of engine; cleanliness; count of defects; size of defects; location of defects; operation of the engine; operation of the exhaust; functionality of the interior electronics; presence of rust; and color.
11. A system for valuing an item, comprising: a client web interface configured to input identification information for the item and a digital representation of the item; and
- a computer system configured to:
- obtain item specifications for the item using the identification information;
 - determine a preliminary condition indication based on the item specifications;
 - determine, from the digital representation using a machine learning (ML) engine, a first item attribute;
 - determine, from the first item attribute, a first condition indication adjustment;
 - determine an adjusted condition indication based on the preliminary condition indication and the first condition indication adjustment;
 - obtain market valuation data of the item based on the identification information; and
 - determine a value of the item based on the market valuation data and the adjusted condition indication.
12. The system of claim 11, wherein the computer system is further configured to:
- determine, from the digital representation using the ML engine, a second item attribute;
 - determine, from the second item attribute, a second condition indication adjustment; and
 - determine the adjusted condition indication based on the preliminary condition indication, the first condition indication adjustment, and the second condition indication adjustment.
13. The system of claim 11, wherein determining, from the digital representation using the ML engine, the first item attribute comprises:
- obtaining a first plurality of digital representations of first similar items;
 - labelling the first plurality of digital representations of first similar items with first training attributes; and
 - training the ML engine with the labelled first plurality of digital representations.
14. The system of claim 11, wherein determining, from the digital representation using the ML engine, the second item attribute comprises:
- obtaining a second plurality of digital representations of second similar items;
 - labelling the second plurality of digital representations of second similar items with second training attributes; and
 - training the ML engine with the labelled second plurality of digital representations.
15. The system of claim 11, wherein the computer system is further configured to:
- obtain history data of the item;
 - wherein determining the preliminary condition indication is further based on the history data.
16. The system of claim 11, wherein the computer system is further configured to:
- obtain ownership information of the item;
 - wherein determining the preliminary condition indication is further based on the ownership information.
17. The system of claim 11, wherein the item is a used vehicle.
18. The method of claim 11, wherein the digital representation comprises a visual representation.
19. The method of claim 11, wherein the digital representation comprises an audio representation.
20. A computer-implemented method of training a neural network for item valuation, comprising:

collecting a set of digital representations of a plurality of similar items;
applying one or more transformations to create a modified set of digital representations of the similar items;
creating an unsupervised learning first training set comprising the modified set of digital representations;
training the neural network using the unsupervised learning first training set;
creating a labeled set of digital representations by labeling a subset of the modified set of digital representations and the collected set of digital representations to identify at least a first attribute and a second attribute;
creating a supervised learning second training set comprising the labeled set of digital representations;
training the neural network using the supervised learning second training set;
inputting a new set of digital representations to the trained neural network;
applying outputs of the trained neural network to a valuation predictor and obtaining valuation results;
collecting a set of rewards and/or punishments from the valuation results; and
further training the neural network, through reinforcement learning, using the set of rewards and/or punishments.

* * * * *