



(19) **United States**

(12) **Patent Application Publication**
Cordell et al.

(10) **Pub. No.: US 2020/0226395 A1**

(43) **Pub. Date: Jul. 16, 2020**

(54) **METHODS AND SYSTEMS FOR DETERMINING WHETHER AN OBJECT IS EMBEDDED IN A TIRE OF A VEHICLE**

B60R 1/04 (2006.01)
B60R 1/062 (2006.01)
G01S 19/48 (2006.01)
G08G 1/14 (2006.01)
H04N 7/18 (2006.01)

(71) Applicant: **Xevo Inc.**, Bellevue, WA (US)

(52) **U.S. Cl.**

(72) Inventors: **John Palmer Cordell**, Los Angeles, CA (US); **Robert Victor Welland**, Seattle, WA (US); **Samuel James McKelvie**, S. Seattle, WA (US); **John Hayes Ludwig**, Bellevue, WA (US)

CPC *G06K 9/00845* (2013.01); *G06N 20/00* (2019.01); *G08G 1/096861* (2013.01); *G08G 1/0969* (2013.01); *H04L 67/10* (2013.01); *G08G 1/096811* (2013.01); *H04W 4/44* (2018.02); *G01C 21/3602* (2013.01); *G08G 1/0116* (2013.01); *B60Q 9/008* (2013.01); *G06K 9/00805* (2013.01); *G06K 9/00979* (2013.01); *G06K 9/66* (2013.01); *G08G 1/166* (2013.01); *G08G 1/0112* (2013.01); *G08G 1/0129* (2013.01); *G08G 1/04* (2013.01); *B60R 1/00* (2013.01); *B60R 1/04* (2013.01); *B60R 1/062* (2013.01); *G06K 9/00281* (2013.01); *G06K 9/00604* (2013.01); *G06K 9/00791* (2013.01); *G01C 21/3605* (2013.01); *G01S 19/48* (2013.01); *G06K 9/00798* (2013.01); *G06K 9/00812* (2013.01); *G06K 9/00832* (2013.01); *G08G 1/096888* (2013.01); *G08G 1/143* (2013.01); *H04N 7/188* (2013.01); *H04W 4/02* (2013.01)

(21) Appl. No.: **16/827,635**

(22) Filed: **Mar. 23, 2020**

Related U.S. Application Data

(63) Continuation of application No. 15/672,747, filed on Aug. 9, 2017.

(60) Provisional application No. 62/372,999, filed on Aug. 10, 2016.

Publication Classification


(51) **Int. Cl.**

G06K 9/00 (2006.01)
H04W 4/02 (2006.01)
G08G 1/0968 (2006.01)
G08G 1/0969 (2006.01)
H04L 29/08 (2006.01)
H04W 4/44 (2006.01)
G01C 21/36 (2006.01)
G08G 1/01 (2006.01)
B60Q 9/00 (2006.01)
G06K 9/66 (2006.01)
G08G 1/16 (2006.01)
G08G 1/04 (2006.01)
B60R 1/00 (2006.01)

(57)

ABSTRACT

A method and/or system is able to improve vehicle safety by determining if an object is embedded in a tire of the vehicle. Audio data is received from a microphone that is positioned to capture sounds of the tire moving on the road. The speed of the vehicle is also obtained, where the speed overlaps the same timeframe of when the sounds of the tire are captured by the microphone. An object is determined to be embedded in the tire based on a frequency analysis of the received audio data relative to the speed of the vehicle. And an alert is output to the driver of the vehicle indicating that the object is embedded in the tire.






FIG. 1A




FIG.1B

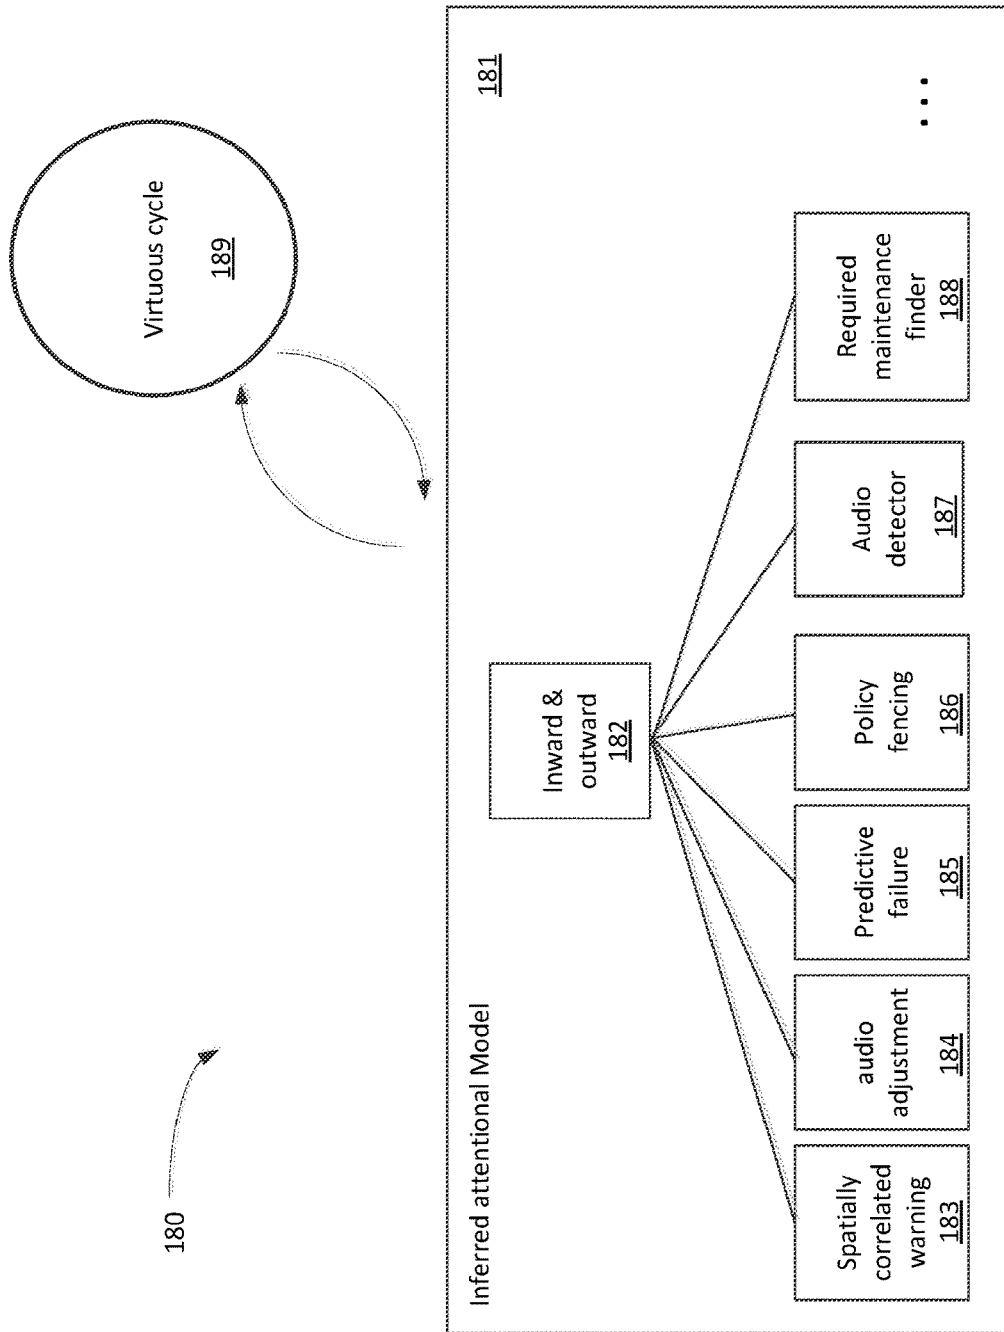


FIG. 1C






FIG. 1D

Using positional metadata to visually inform the driver



Crash when driver has scanned forward appropriately

Crash when driver has turned right but failed to check in the right before the turn



Mirrors when driver has been checking them at appropriate intervals

Mirrors when driver has failed to check the right mirror at appropriate intervals

FIG. 1E




FIG. 1F




FIG. 1G

i
l
c

1

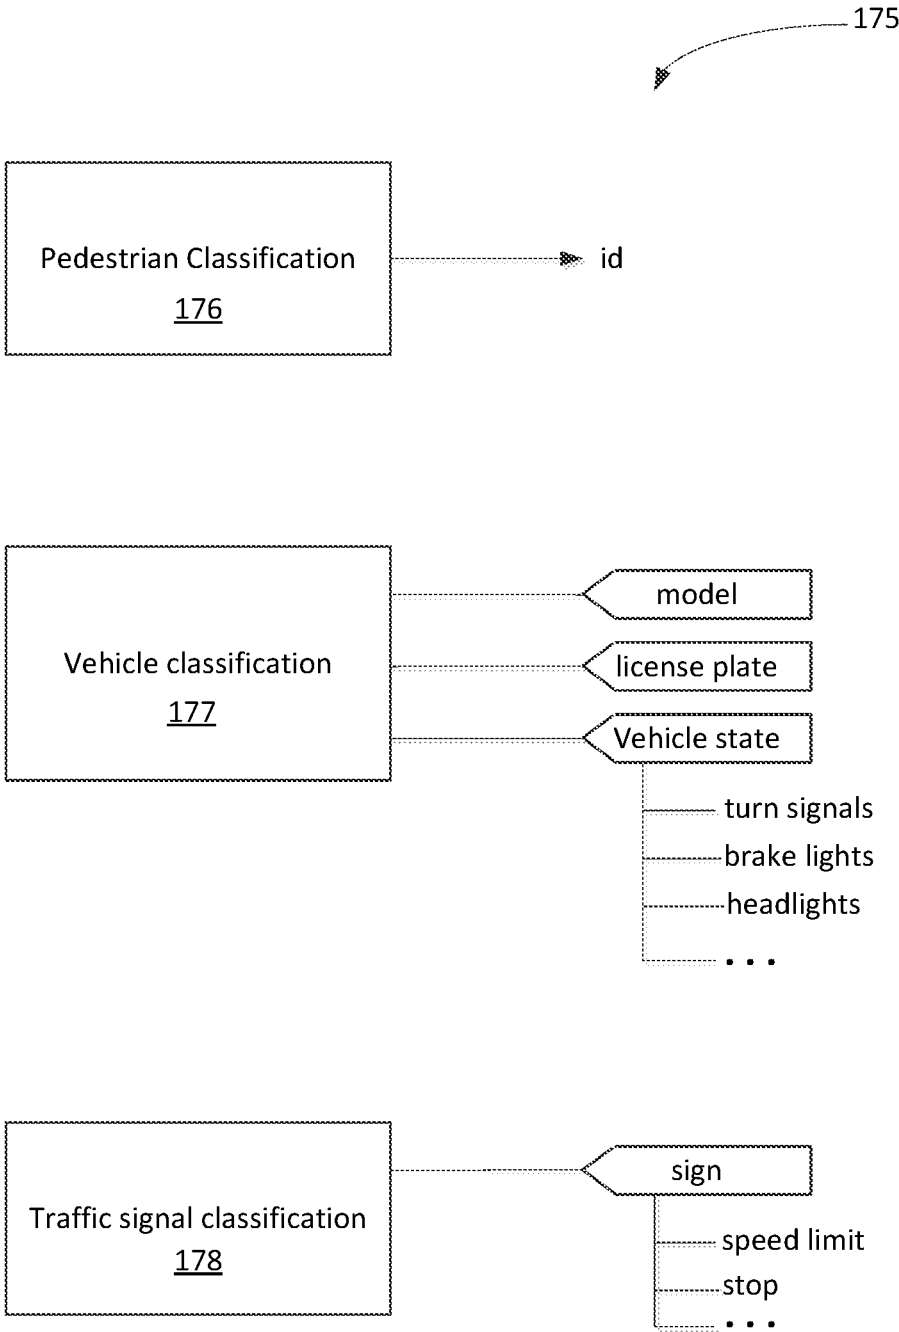


FIG. 1H

200




FIG. 2A

230




FIG. 2B

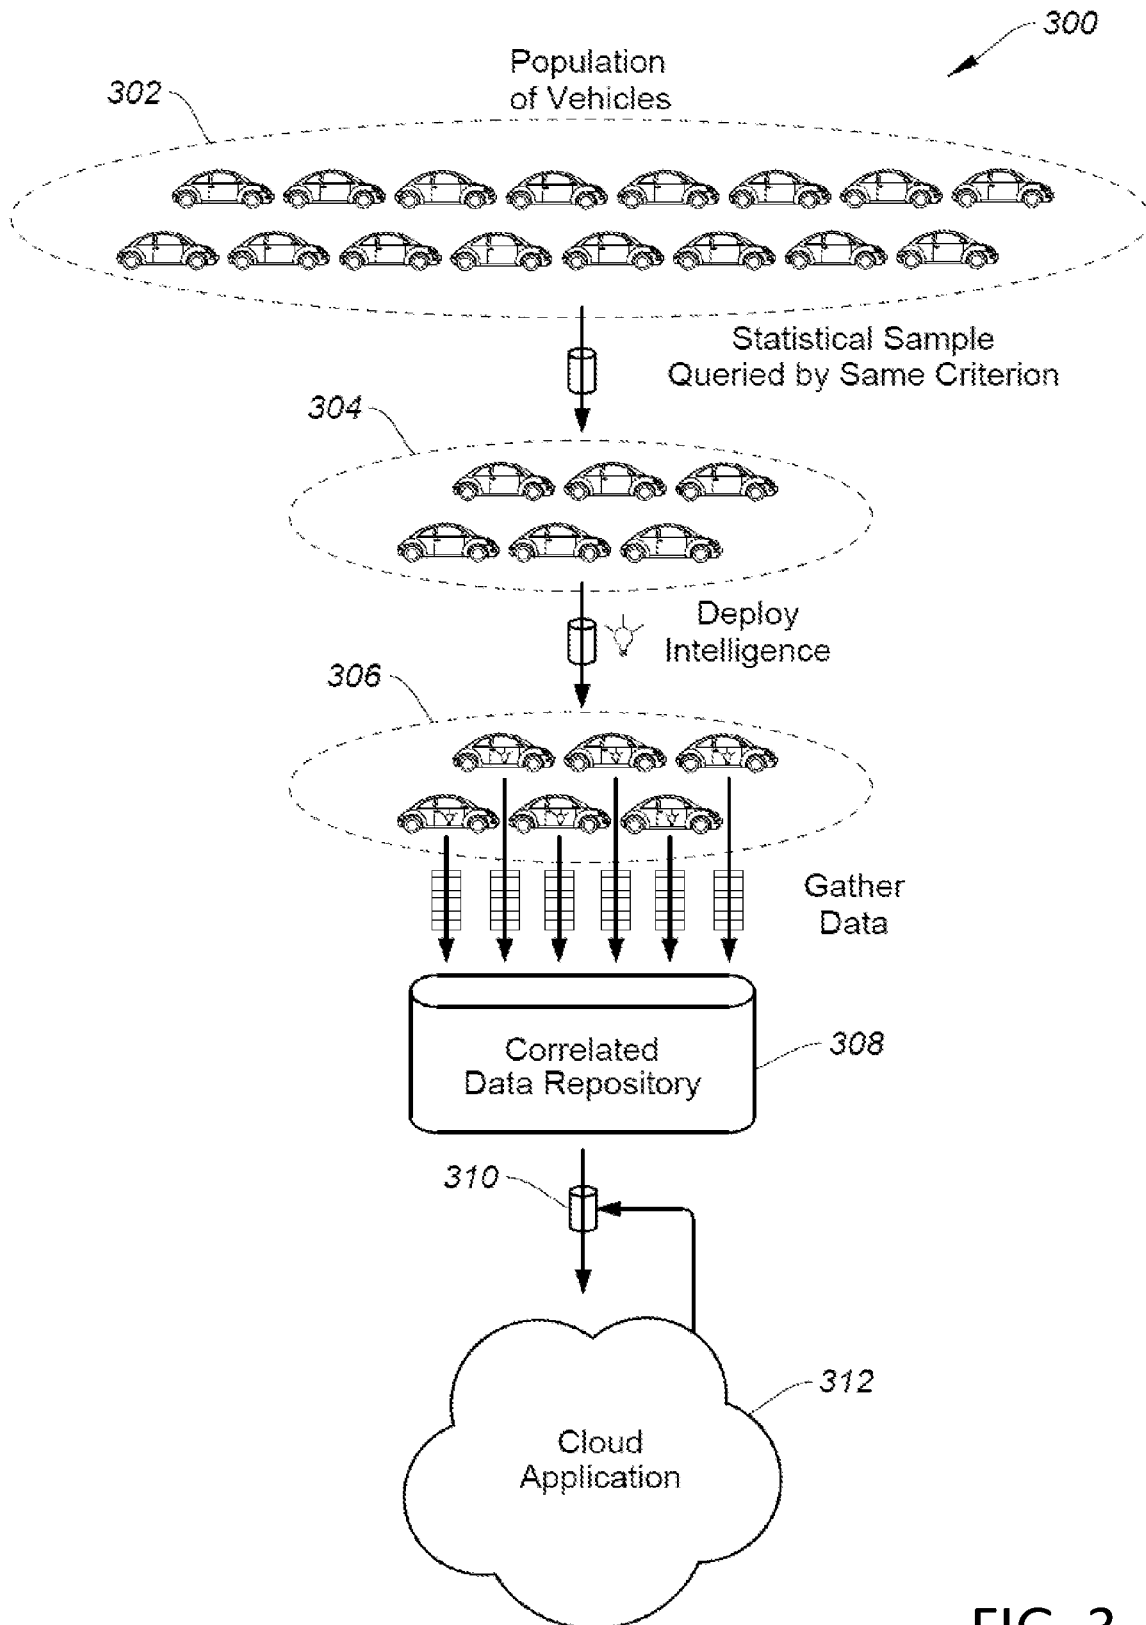


FIG. 3

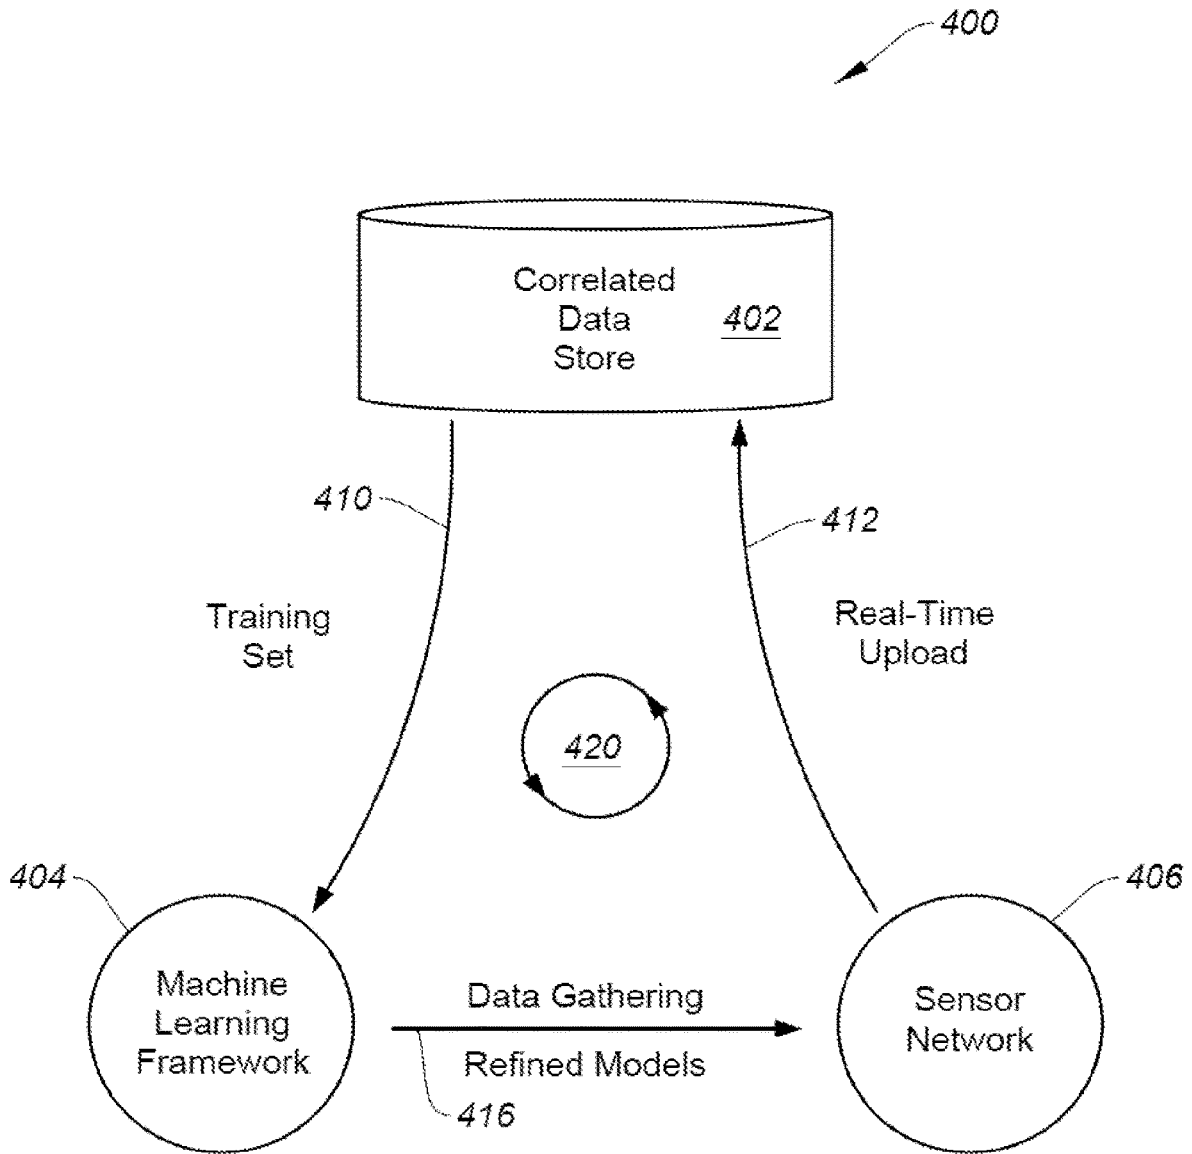


FIG. 4

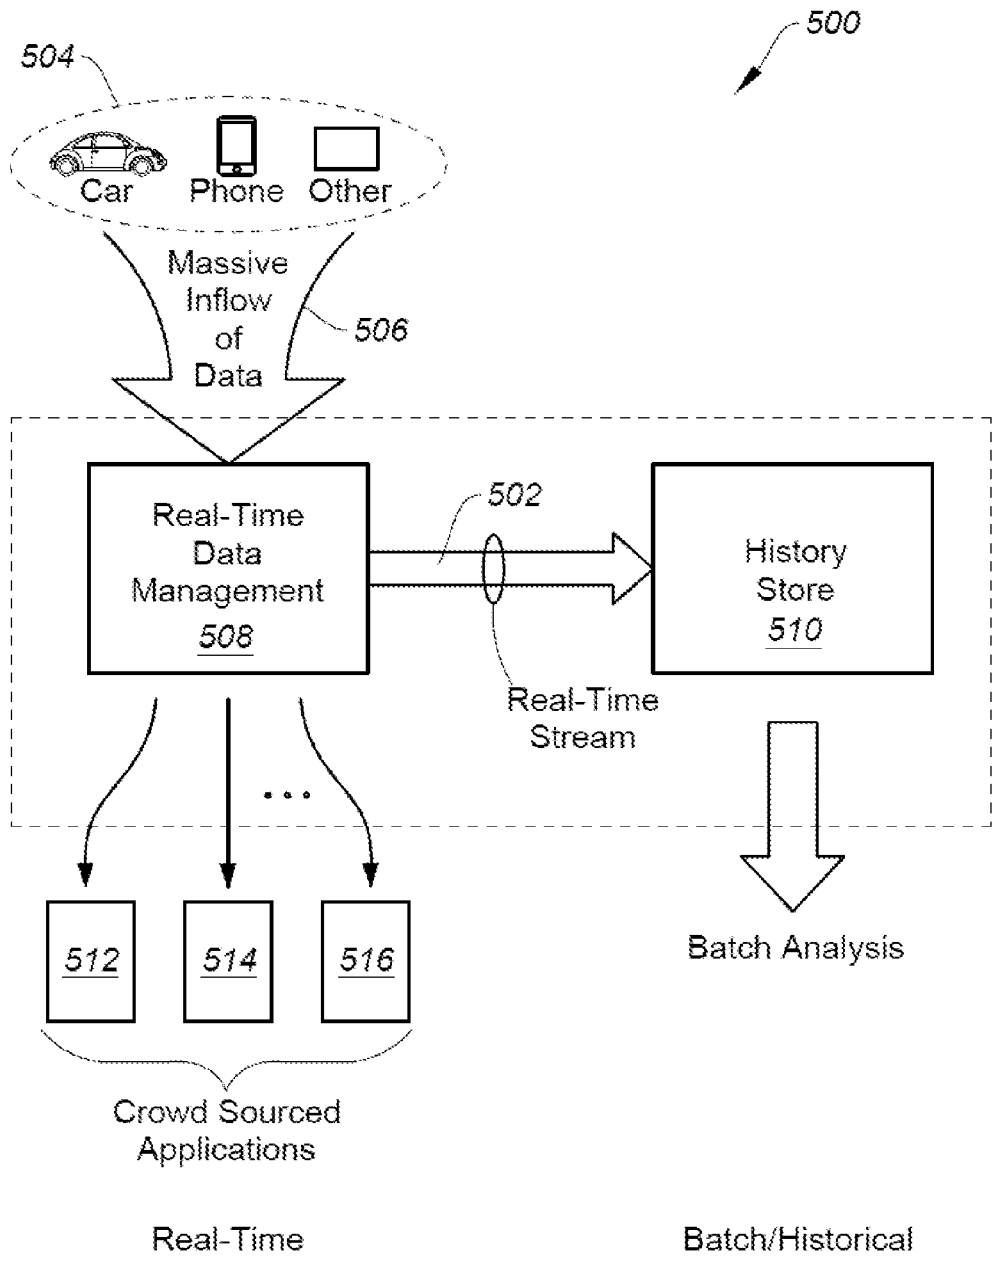


FIG. 5




FIG. 6





FIG. 7






FIG. 9




FIG. 10

1100

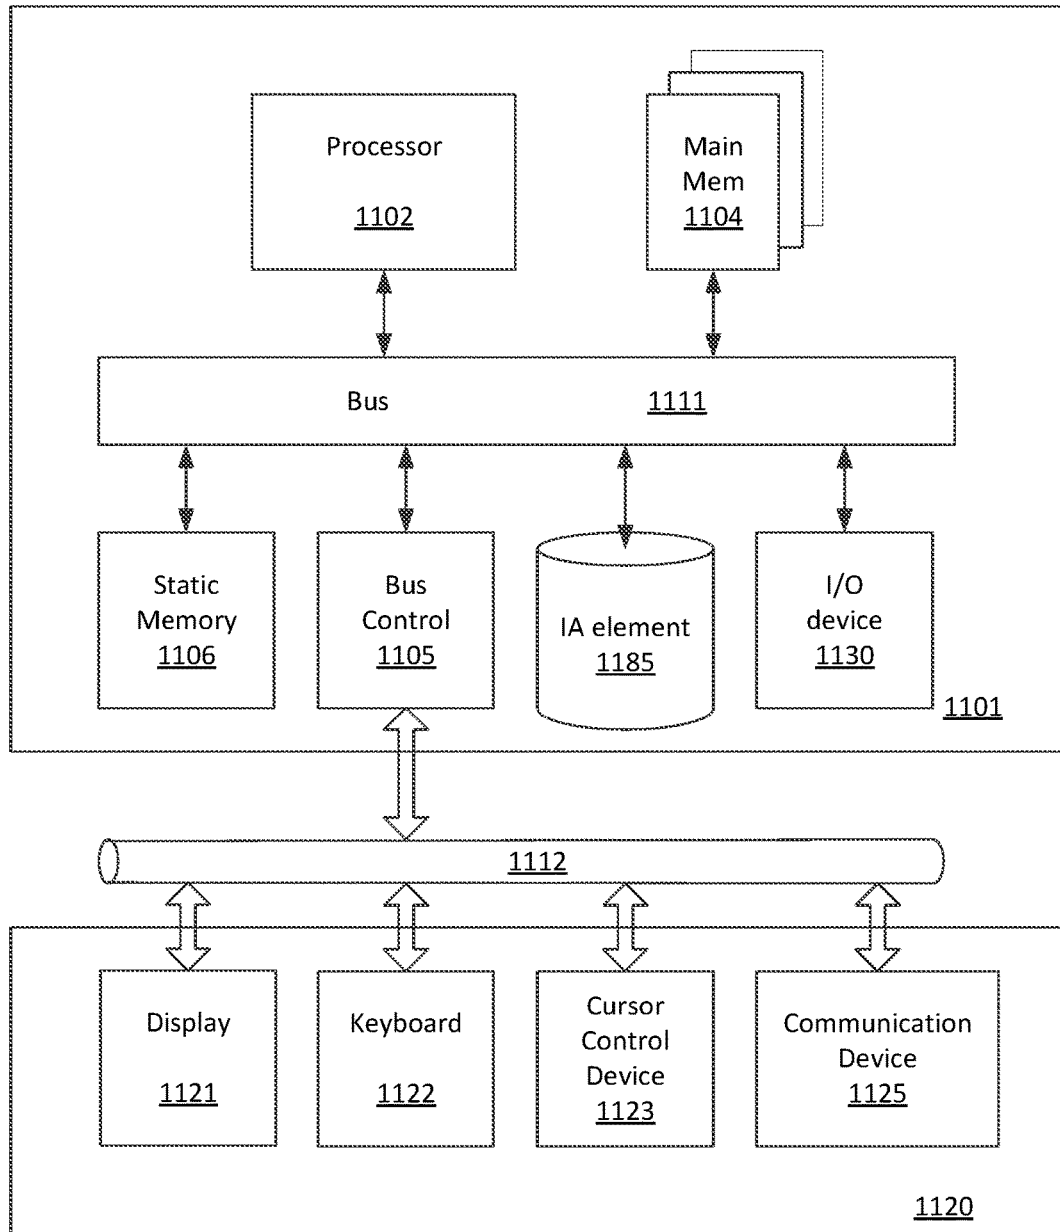


FIG 11

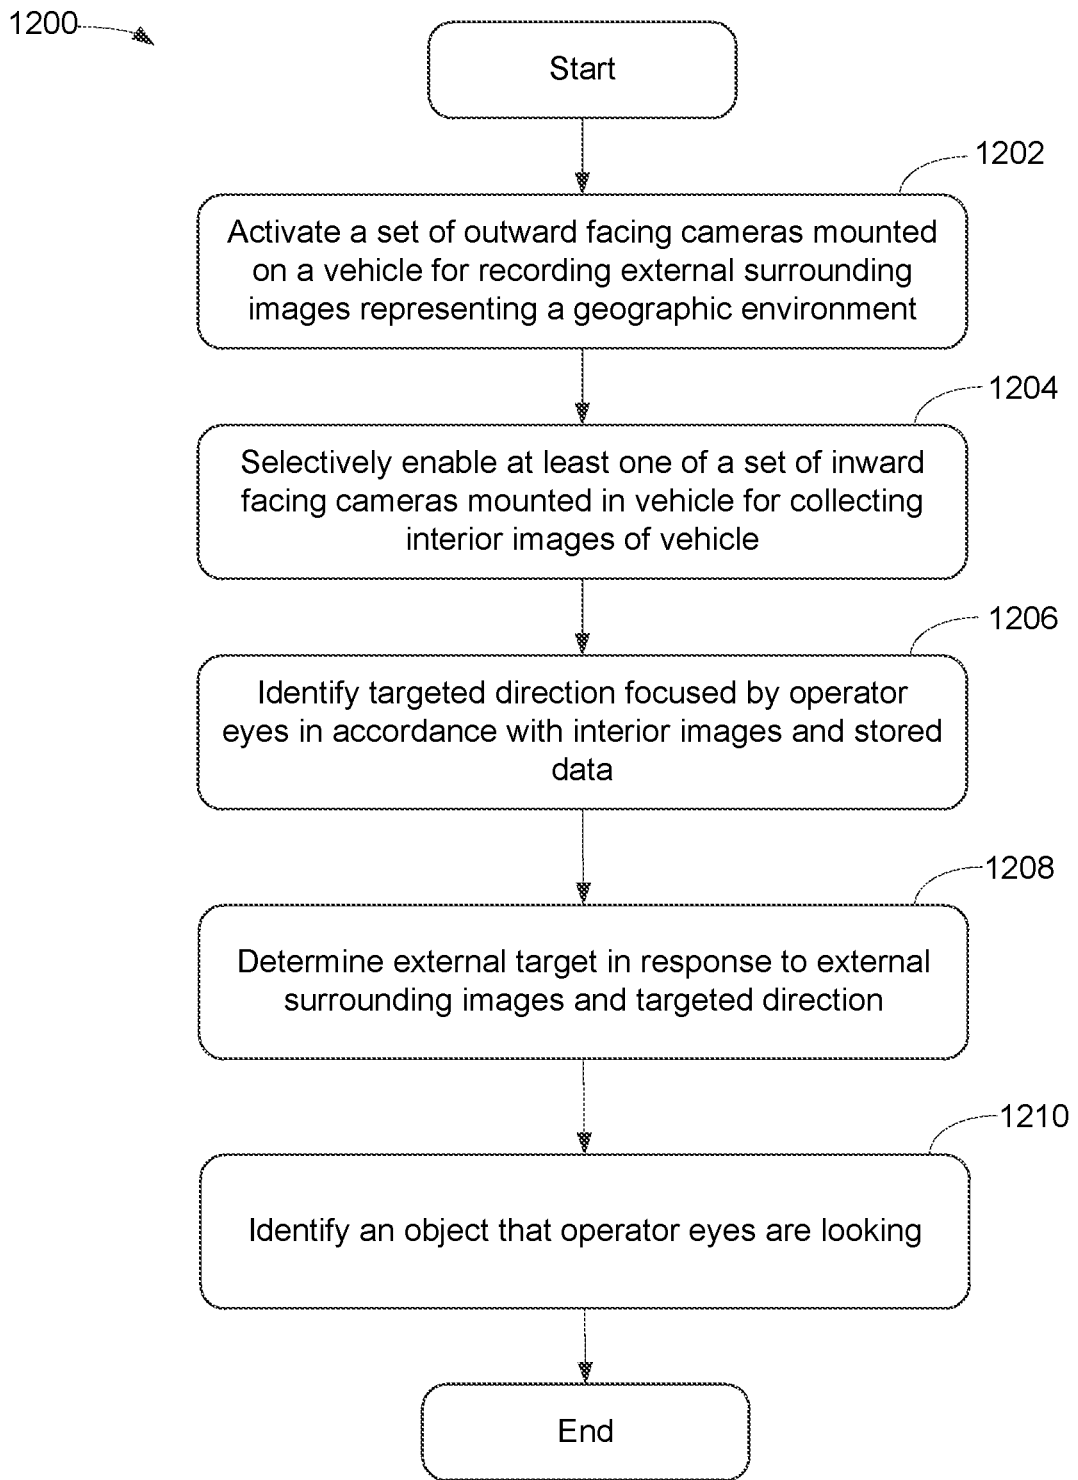


FIG 12

**METHODS AND SYSTEMS FOR
DETERMINING WHETHER AN OBJECT IS
EMBEDDED IN A TIRE OF A VEHICLE**

PRIORITY

[0001] This application claims the benefit of priority based upon U.S. Provisional Patent Application having an application Ser. No. 62/372,999, filed on Aug. 10, 2016 and having a title of “Method and System for Providing Information Using Collected and Stored Metadata,” as well as U.S. Non-provisional Patent Application having an application Ser. No. 15/672,747 and having a title of “Method and Apparatus for Providing Information Via Collected and Stored Metadata Using Inferred Attentional Model,” which are hereby incorporated by reference in its entirety.

FIELD

[0002] The exemplary embodiment(s) of the present invention relates to the field of communication networks. More specifically, the exemplary embodiment(s) of the present invention relates to operating an intelligent machine using a virtuous cycle between cloud, machine learning, and containerized sensors.

BACKGROUND

[0003] With increasing popularity of automation and intelligent electronic devices, such as computerized machines, IoT (the Internet of Things), smart vehicles, smart phones, drones, mobile devices, airplanes, artificial intelligence (“AI”), the demand of intelligent machines and faster real-time response are increasing. To properly provide machine learning, a significant number of pieces, such as data management, model training, and data collection, needs to be improved.

[0004] A conventional type of machine learning is, in itself, an exploratory process which may involve trying different kinds of models, such as convolutional, RNN (recurrent neural network), et cetera. Machine learning or training typically concerns a wide variety of hyper-parameters that change the shape of the model and training characteristics. Model training generally requires intensive computation. As such, real-time response via machine learning model can be challenging.

[0005] A drawback associated with traditional automobile or vehicle is that a vehicle typically makes some decisions with limited knowledge of the context or environment in which it operates. Also, a vehicle has limited ability to participate in creation of user or operator experience.

SUMMARY

[0006] One embodiment of the presently claimed invention discloses a method and/or inferred attentional system (“IAS”) capable of enhancing vehicle safety via metadata extraction by an IA model trained by a virtuous cycle containing sensors, machine learning center (“MLC”), and cloud based network (“CBN”). In one aspect, IAS includes a set of outward facing camera, inward facing cameras, and vehicle onboard computer (“VOC”). The outward facing cameras collect external images representing a surrounding environment in which the vehicle operates. The collecting external images include obtaining real-time images relating to at least one of road, buildings, traffic lights, pedestrian, and retailers.

[0007] The inward facing cameras collect internal images including operator facial expression representing at least operator’s attention. The collecting internal images include a set of interior sensors capable of obtaining data relating to at least one of operator’s eyes, facial expression, driver, and passage. The obtaining data relating to at least one of operator’s eyes, facial expression, driver, and passage further includes at least one camera is capable of detecting direction of where the operator is looking.

[0008] The VOC is configured to identify operator’s attention in response to the collected internal images and the collected external images. In one aspect, the VOC is wirelessly connected to Internet capable of communicating with cloud data such as the virtuous cycle for modifying the IA model. The VOC includes a pipeline processor capable of identify what the operator is looking based on the direction in which the operator is looking based on collected by interior cameras and exterior objects captured by the exterior cameras. In addition, the VOC has a warning component capable of issuing a warning sound based on exterior data obtained by exterior cameras and interior data collected by interior cameras.

[0009] The system or IAS also includes a set of audio sensors that are coupled to the VOC and configured to provide metadata relating to audio data. The audio sensors include exterior audio sensors collecting exterior sound outside of the vehicle and interior audio sensors collecting interior sound inside of the vehicle.

[0010] In one embodiment of the presently claimed invention discloses a method for assisting vehicle operation via metadata extraction processed by the IA. The method is capable of activating a set of outward facing cameras mounted on a vehicle for recording external surrounding images representing a geographic environment in which the vehicle operates. In addition, at least one of a set of inward facing cameras which are mounted in the vehicle is selectively enabled for collecting interior images of the vehicle. Upon identifying targeted direction focused by operator eyes in accordance with the interior images and stored data managed by the IA model, the external target is determined in response to the external surrounding images and the targeted direction. An object that the operator eyes are looking is identified based on the external target and information supplied by the IA model.

[0011] After tracking or monitoring surrounding environmental event in accordance with the external surrounding images and information provided by the IA model trained by a virtuous cycle, a potential vehicle collision is identified in response to the surrounding environmental event. In one embodiment, the method is able to provide a spatially correlated audio warning to the operator if the potential vehicle collision is determined. The method also provides automatic adjustment of interior entertainment volume when the spatially correlated audio warning is activated. Upon providing a spatially correlated visual warning to the operator if the potential vehicle collision is determined, the process is capable of adjusting moving direction of the vehicle to avoid potential collision in accordance with a set of predefined collision avoidance policy facilitated by the IA model which is trained by a virtuous cycle via a cloud computing.

[0012] In an alternative embodiment, the presently claimed invention discloses a process of facilitating a prediction of imminent machine failure via metadata generated

by embedded sensors. The process, for example, is able to activate at least one embedded sensor attached to a mechanical component in a vehicle for collecting metadata from the mechanical component. After processing the metadata by a local pipelined processing unit to convert metadata format to a transferrable metadata format, the process uploads the converted metadata to cloud via a wireless communication network for training a failure prediction model including analyzing the converted metadata according to a plurality of failure samples relating to the mechanical component. A failure predication may be pushed to the vehicle to indicate an imminent machine failure. Upon activating a set of outward facing video cameras for aggregating a predefined set of images relating to external geographic environment in which the vehicle operates, the aggregated images is uploaded to a cloud via a wireless communications network for detecting required maintenance.

[0013] One embodiment of the presently claimed invention is capable of facilitating identification of attentional state. Upon presenting to a first software application layer configured to build a first real-time object model representing a first contextual state of exterior of a vehicle, a second software application layer configured to build a second real-time object model representing a second contextual state of interior of the vehicle is presented. After generating a third software application to build a third real-time object model representing attentional state of occupants of the vehicle, the attentional state is determined in accordance with relationships between the first and second contextual states and the attentional state of occupants.

[0014] Additional features and benefits of the exemplary embodiment(s) of the present invention will become apparent from the detailed description, figures and claims set forth below.

BRIEF DESCRIPTION OF THE DRAWINGS

[0015] The exemplary embodiment(s) of the present invention will be understood more fully from the detailed description given below and from the accompanying drawings of various embodiments of the invention, which, however, should not be taken to limit the invention to the specific embodiments, but are for explanation and understanding only.

[0016] FIGS. 1A-1B are block diagrams illustrating a virtuous cycle facilitating an inferred attentional system (“IAS”) which is capable of identifying operator attention via IA model trained by the virtuous cycle in accordance with one embodiment of the present invention;

[0017] FIGS. 1C-1E are block diagrams illustrating an IA model containing various components using inward and outward facing cameras via a virtuous cycle in accordance with one embodiment of the present invention;

[0018] FIGS. 1F-111 is a block diagram illustrating a pipeline process of outward facing camera capable of identifying and classifying detected object(s) using a virtuous cycle in accordance with one embodiment of the present invention;

[0019] FIGS. 2A-2B are block diagrams illustrating a virtuous cycle capable of facilitating IA model detection in accordance with one embodiment of the present invention;

[0020] FIG. 3 is a block diagram illustrating a cloud based network using crowdsourcing approach to improve IA model (s) in accordance with one embodiment of the present invention;

[0021] FIG. 4 is a block diagram illustrating an IA model or system using the virtuous cycle in accordance with one embodiment of the present invention;

[0022] FIG. 5 is a block diagram illustrating an exemplary process of correlating IA data in accordance with one embodiment of the present invention;

[0023] FIG. 6 is a block diagram illustrating an exemplary process of real-time data management for IA model in accordance with one embodiment of the present invention;

[0024] FIG. 7 is a block diagram illustrating a crowd sourced application model for IA model in accordance with one embodiment of the present invention;

[0025] FIG. 8 is a block diagram illustrating a method of storing IA related data using a geo-spatial objective storage in accordance with one embodiment of the present invention;

[0026] FIG. 9 is a block diagram illustrating an exemplary approach of analysis engine analyzing IA data in accordance with one embodiment of the present invention;

[0027] FIG. 10 is a block diagram illustrating an exemplary containerized sensor network used for sensing IA related information in accordance with one embodiment of the present invention;

[0028] FIG. 11 is a block diagram illustrating a processing device or computer system which can be installed in a vehicle for facilitating the virtuous cycle in accordance with one embodiment of the present invention; and

[0029] FIG. 12 is a flowchart illustrating a process of IA model or system capable of identifying operator attention in accordance with one embodiment of the present invention.

DETAILED DESCRIPTION

[0030] Embodiments of the present invention are described herein with context of a method and/or apparatus for facilitating detection of operator attention via an inferred attentional system (“IAS”) using an IA model continuously trained by a virtuous cycle containing cloud based network, containerized sensing device, and machine learning center (“MLC”).

[0031] The purpose of the following detailed description is to provide an understanding of one or more embodiments of the present invention. Those of ordinary skills in the art will realize that the following detailed description is illustrative only and is not intended to be in any way limiting. Other embodiments will readily suggest themselves to such skilled persons having the benefit of this disclosure and/or description.

[0032] In the interest of clarity, not all of the routine features of the implementations described herein are shown and described. It will, of course, be understood that in the development of any such actual implementation, numerous implementation-specific decisions may be made in order to achieve the developer’s specific goals, such as compliance with application- and business-related constraints, and that these specific goals will vary from one implementation to another and from one developer to another. Moreover, it will be understood that such a development effort might be complex and time-consuming, but would nevertheless be a routine undertaking of engineering for those of ordinary skills in the art having the benefit of embodiment(s) of this disclosure.

[0033] Various embodiments of the present invention illustrated in the drawings may not be drawn to scale. Rather, the dimensions of the various features may be

expanded or reduced for clarity. In addition, some of the drawings may be simplified for clarity. Thus, the drawings may not depict all of the components of a given apparatus (e.g., device) or method. The same reference indicators will be used throughout the drawings and the following detailed description to refer to the same or like parts.

[0034] In accordance with the embodiment(s) of present invention, the components, process steps, and/or data structures described herein may be implemented using various types of operating systems, computing platforms, computer programs, and/or general purpose machines. In addition, those of ordinary skills in the art will recognize that devices of a less general purpose nature, such as hardware devices, field programmable gate arrays (FPGAs), application specific integrated circuits (ASICs), or the like, may also be used without departing from the scope and spirit of the inventive concepts disclosed herein. Where a method comprising a series of process steps is implemented by a computer or a machine and those process steps can be stored as a series of instructions readable by the machine, they may be stored on a tangible medium such as a computer memory device (e.g., ROM (Read Only Memory), PROM (Programmable Read Only Memory), EEPROM (Electrically Erasable Programmable Read Only Memory), FLASH Memory, Jump Drive, and the like), magnetic storage medium (e.g., tape, magnetic disk drive, and the like), optical storage medium (e.g., CD-ROM, DVD-ROM, paper card and paper tape, and the like) and other known types of program memory.

[0035] The term “system” or “device” is used generically herein to describe any number of components, elements, sub-systems, devices, packet switch elements, packet switches, access switches, routers, networks, computer and/or communication devices or mechanisms, or combinations of components thereof. The term “computer” includes a processor, memory, and buses capable of executing instruction wherein the computer refers to one or a cluster of computers, personal computers, workstations, mainframes, or combinations of computers thereof.

[0036] One embodiment of the presently claimed invention discloses an IAS capable of enhancing vehicle safety via metadata extraction by an IA model trained by a virtuous cycle containing sensors, MLC, and cloud based network (“CBN”). In one aspect, the system or IAS includes a set of outward facing camera, inward facing cameras, and vehicle onboard computer (“VOC”). The outward facing cameras collect external images representing a surrounding environment in which the vehicle operates. The inward facing cameras collect internal images including operator facial expression representing at least operator’s attention. The VOC is configured to identify operator’s attention in response to the collected internal images and the collected external images.

[0037] In an alternative embodiment, the presently claimed invention discloses an IA process of facilitating a prediction of imminent machine failure via metadata generated by embedded sensors. The process, for example, is able to activate at least one embedded sensor attached to a mechanical component in a vehicle for collecting metadata from the mechanical component. After processing the metadata by a local pipelined processing unit to convert metadata format to a transferrable metadata format, the process uploads the converted metadata to cloud via a wireless communication network for training a failure prediction

model including analyzing the converted metadata according to a plurality of failure samples relating to the mechanical component. A failure predication may be pushed to the vehicle to indicate an imminent machine failure.

[0038] In addition, upon activating a set of outward facing video cameras for aggregating a predefined set of images relating to external geographic environment in which the vehicle operates, the aggregated images is uploaded to a cloud via a wireless communications network for detecting required maintenance.

[0039] FIG. 1A is a block diagram **100** illustrating a virtuous cycle facilitating an IAS which is capable of identifying operator attention via IA model trained by the virtuous cycle in accordance with one embodiment of the present invention. Diagram **100** illustrates a virtuous cycle containing a vehicle **102**, CBN **104**, and MLC **106**. In one aspect, MCL **106** can be located remotely or in the cloud. Alternatively, MCL **106** can be a part of CBN **104**. It should be noted that the underlying concept of the exemplary embodiment(s) of the present invention would not change if one or more blocks (circuit or elements) were added to or removed from diagram **100**.

[0040] Vehicle **102**, in one example, can be a car, automobile, bus, train, drone, airplane, truck, and the like, and is capable of moving geographically from point A to point B. To simplify forgoing discussing, the term “vehicle” or “car” is used. Vehicle **102** includes wheels with ABS (anti-lock braking system), body, steering wheel **108**, exterior or outward facing cameras **125**, interior (or 360° (degree)) or inward facing camera **126**, antenna **124**, onboard controller or VOC **123**, and operator (or driver) **109**. It should be noted that outward facing cameras and/or inward facing cameras **125-126** can be installed at front, side-facing, stereo, and inside of vehicle **102**. In one example, vehicle **102** also includes various sensors which senses information related to vehicle state, vehicle status, driver actions, For example, the sensors, not shown in FIG. 1A, are able to collect information, such as audio, ABS, steering, braking, acceleration, traction control, windshield wipers, GPS (global positioning system), radar, ultrasound, lidar (Light Detection and Ranging), and the like.

[0041] VOC or onboard controller **123** includes CPU (central processing unit), GPU (graphic processing unit), memory, and disk responsible for gathering data from outward facing or exterior cameras **125**, inward facing or interior cameras **126**, audio sensor, ABS, traction control, steering wheel, CAN-bus sensors, and the like. In one aspect, VOC **123** executes IA model received from MLC **106**, and interfaces with antenna **124** to communicate with CBN **104** via a wireless communication network **110**. Note that wireless communication network includes, but not limited to, WIFI, cellular network, Bluetooth network, satellite network, or the like. A function of VOC **123** is to gather or capture real-time surrounding information as well as exterior information when vehicle **102** is moving.

[0042] CBN **104** includes various digital computing systems, such as, but not limited to, server farm **120**, routers/switches **121**, cloud administrators **119**, connected computing devices **116-117**, and network elements **118**. A function of CBN **104** is to provide cloud computing which can be viewed as on-demand Internet based computing service with enormous computing power and resources. Another function of CBN **104** is to improve or refine IA labeled data via correlating captured real-time data with relevant cloud data.

The refined IA labeled data is subsequently passed to MLC 106 for model training via a connection 112.

[0043] MLC 106, in one embodiment, provides, refines, trains, and/or distributes models 115 such as IA model based on information or data such as IA labeled data provided from CBN 104. It should be noted that the machine learning makes predictions based on models generated and maintained by various computational algorithms using historical data as well as current data. A function of MLC 106 is that it is capable of pushing information such as revised IA model to vehicle 102 via a wireless communications network 114 in real-time.

[0044] To identify or collect operator attention of vehicle 102, an onboard IA model which could reside inside of VOC 123 receives a triggering event or events from built-in sensors such as ABS, wheel slippery, turning status, engine status, and the like. The triggering event or events may include, but not limited to, activation of ABS, rapid steering, rapid breaking, excessive wheel slip, activation of emergency stop, and on. Upon receiving triggering events via vehicular status signals, the recording or recorded images captured by inward facing camera or 360 camera are reworded from an earlier time stamp leading to the receipt of triggering event(s) for identifying IA labeled data which contains images considered to be dangerous driving. After correlation of IA labeled data with historical sampling data at CBN, the IA model is retrained and refined at MLC 106. The retrained IA model is subsequently pushed back onto vehicle 102.

[0045] In one embodiment, triggering events indicate an inattentive or distracted driver. For example, upon detecting a potential dangerous event, CBN 104 issues warning signal to driver or operator 109 via, for instance, a haptic signal, or shock to operator 109 notifying a potential collision. In addition, the dangerous event or events are recorded for report. It should be noted that a report describing driver's behavior as well as number occurrence relating to dangerous events can be useful. For example, such report can be obtained by insurance company for insurance auditing, by law enforcement for accident prevention, by city engineers for traffic logistics, or by medical staff for patient safety.

[0046] During an operation, inward facing camera 126 captures facial images of driver or operator 109 including the location in which operator's eyes focusing. Upon verifying with CBN 104, a focal direction 107 of operator 109 is identified. After obtaining and processing external images relating to focal direction 107, a possible trajectory 105 in which the location is looked at is obtained. Trajectory 105 and focal direction 107 are subsequently processed and combined in accordance with stored data in the cloud. The object, which is being looked at by operator 109, is identified. In this example, the object is a house 103 nearby the road.

[0047] An advantage of using an IAS is to reduce traffic accidents and enhance public safety. With employment of IAS, the full context of vehicle, both inside and out, is a very rich set of information. It encompasses simple things like the current value of the multitude of sensors in a vehicle, such as pedal position, steering wheel position, mirror setting, seat setting, engine RPM, whether the seat belts are clipped in, internal and external temperature, et cetera. With the advent of machine learning, a broad class of derived data and metadata can be extracted from sensors and be used to improve the user experience of being in or driving a vehicle.

It should be noted that the extracted data includes confidence and probability metrics for each data element that the machine learning models emit. Such data, which changes in real-time, is presented to an application layer that can use the full context of vehicle in real-time.

[0048] FIG. 1B illustrates a block diagram 140 showing an operator or driver monitored by IAS able to identify operator attention via a virtuous cycle in accordance with one embodiment of the present invention. Diagram 140 illustrates a driver 148, inward facing camera(s) 142, and exterior camera 144. In one aspect, camera 142, also known as interior camera or 360 camera, monitors or captures driver's facial expression 146 and/or driver (or operator) body language. Upon reading IA status 149 which indicates stable with accelerometer, ahead with gaze, hands on steering wheel (no texting), the IAS concludes that driver is behaving normally.

[0049] In one embodiment, IAS is able to detect which direction driver 148 is looking, whether driver 148 is distracted, whether driver 148 is texting, whether identity of driver is determined via a facial recognition process, and/or where driver 148 pays attention. It should be noted that the car may contain multiple forward facing cameras (or 360-degree camera(s)) 144 capable of capturing a 360 view which can be used to correlate with other views to identify whether driver 148 checks back-view mirror to see cars behind the vehicle or checks at side view of vehicle when the car turns. Based on observed IA, the labeled data showing looking at the correct spots based on traveling route of car can illustrate where the driver pays attention. Alternatively, the collected images or labeled data can be used to retrain IA model which may predict the safety rating for driver 148.

[0050] During an operation, the interior images captured by inward facing camera(s) 142 can show a location in which operator 148 is focusing based on relative eye positions of operator 148. Once the direction of location such as direction 145 is identified, IAS obtains external images captured by outward facing camera(s) 144. After identifying image 145 is where operator pays attention based on direction 145, the image 145 is recorded and process. Alternatively, if IAS expects operator 148 should look at the direction 145 based on current traveling speed, whether condition, visibility, and traffic condition, operator 148 actually is looking at a house 141 based in trajectory view 143 based on captured images, a warning signal will be activated.

[0051] It should be noted that the labeled data should include various safety parameters such as whether the driver looks left and/or right before crossing an intersection and/or whether the driver gazes at correct locations while driving. IAS also includes multiple sensor or sensors, Lidar, radar, sonar, thermometers, audio detector, pressure sensor, air-flow, optical sensor, infrared reader, speed sensor, altitude sensor, and the like. The information related to IA can change based on occupant(s) behavior in the vehicle or car. For example, if occupants are noisy, loud radio, shouting, drinking, eating, dancing, the occupants behavior can affect overall IA contributes to bad driving behavior.

[0052] FIG. 1C is a block diagram 180 illustrating an IA model containing various components using inward and outward facing cameras via a virtuous cycle in accordance with one embodiment of the present invention. Diagram 180 includes an IAS containing IA model 181 and virtuous cycle 189. IAS which includes IA model 181 may be situated in

vehicle, virtuous cycle **189**, or both. It should be noted that the underlying concept of the exemplary embodiment(s) of the present invention would not change if one or more blocks (components or elements) were added to or removed from diagram **180**.

[0053] IA model **181** containing inward & outward sensors **182**, in one aspect, includes multiple components, such as, but not limited to, spatially correlated warning component **183**, audio adjustment **184**, predictive failure component **185**, policy fencing component **186**, audio detector **187**, and/or required maintenance finder **188**. While spatially correlated warning component **183** is able to provide warnings using emulated real-time spatial effect, audio adjustment **184** is configured to automatic adjust in-car entertainment volume based on external events detected. Predictive failure component **185**, in one aspect, uses multiple embedded microphones in or around a machine or device to predict imminent machine failure based on collected and stored data. Policy fencing component **186** is a set of optional rules used to limit certain behaviors of vehicle performance based on detected and stored data. Audio detector **187** is configured to monitor and detect nails and/or screw in a tire. Required maintenance finder **188** is configured to identify potential required maintenance based on outward facing sensors.

[0054] Inward and Outward Facing Video Metadata

[0055] IAS includes an inward and outward sensors containing one or more video cameras and/or input sensors for extracting information relating to what is happening in the field of view of a human operator (such as a vehicle driver, pedestrian, machine operator, etc.). In addition, video cameras facing the operator extract metadata about the operator, such as head pose, gaze direction, activities such as texting, talking on the phone, interacting with the audio system, looking at mobile device, et cetera. By combining extracted metadata streams (inward and outward facing), IA model of where the operator has focused his/her attention can be constructed in real-time. In one example, IA model can be used to devise a “smart” alert or warning system. For example, when a driver is looking at their mobile phone and the upcoming traffic light changes from green to yellow, a warning signal should be issued.

[0056] IA model, in one embodiment, provides data and metadata flowing through the system or IAS. Software components can create a real-time “model” of where the operator’s attention is currently focused. The system or IAS can know where operator eyes are currently gazing at, whether the eyes open or closed, a history of when they have blinked and how long each blink lasted. IA model can also include information from other inputs to the pipeline computing. For example, an audio microphone can detect whether the operator is talking or using a touchscreen device. The model of the user’s attention or IA model can also be used by another software component to devise a user experience, including warnings and alerts that utilizes this understanding of the current user state to optimize and improve the user experience and performance of the system.

[0057] IAS, in one embodiment, includes a surrounding environment event and situation model, not shown in FIG. 1C, is a real-time software model which keeps track of where the user’s attention is focused. A model of what’s happening externally is also constructed using the extracted metadata packets. For example, external events may indicate that “a pedestrian is at location x, y at time t” or “a light that

is 120 feet away changes from green to yellow” or “two people are currently in the room” or “it’s raining.” It should be noted that such metadata events can be extracted and processed in real-time.

[0058] IAS also includes a software component to handle user experience and alerts which uses both IA model and exterior situational model. The component, in one example, is able to use the higher level understanding of what is happening in order to respond appropriately to the current situation. For example, in the case of a user driving a car, if the system knows that the user was looking forward when the light changes to red, there is no need to issue a warning. If, however, the system knows that the user was looking down at their mobile phone when the light changes, it issues an audible warning to the driver about traffic lights.

[0059] Spatially Correlated Audio Warnings

[0060] IAS, in one aspect, includes a spatially correlated warnings component using multi-speaker system to generate warning tones or messages that are specially correlated to the direction from which the threat or alert emanates. Using a multi-camera system to extract real-time metadata regarding what is happening outside of a vehicle, the location of the activity that is generating the alert can be used to localize the audio alert within the vehicle. For example, if a bicycle is about collide with the right front side of a vehicle, the warning tone can be placed in a sound-field such that it appears to the operator to be coming from the right front side. A benefit of using a spatially correlated warning system is to direct an operator or driver that intuitively focuses his/her attention quickly in the direction that is important.

[0061] To manage level control of individual speakers in a multi-speaker system, the extracted high-level metadata includes video, audio, and other sensors, as well as a software component integrating that information. The software component issuing warnings or alerts should have information relating to the location and direction of the event that is triggering the alert. In a stereo or quadraphonic speaker system, the volume of each individual speaker can be independently adjusted. When the same signal is sent at various volumes to a set of speakers, the result is to “place” the apparent location of the source of the sound to the listener. The system calculates the appropriate relative attenuation values for each speaker that will result in a warning or alert tone or sound emanating from the same direction that the danger or situation is originated.

[0062] FIG. 1D is a block diagram **190** illustrating a process of spatially correlated warnings component via inward and outward facing cameras using a virtuous cycle in accordance with one embodiment of the present invention. Diagram **190** includes a vehicle **193** and source of ambulance siren **192**. Vehicle **193** further includes four (4) speakers **194-197**. IAS, in one embodiment, is able to calculate the geo-location of siren origin and emulating the siren inside of car **193** based on surrounding environment **191** of car **193**. For example, once the surrounding environment is established based on extracted data and stored data, IAS is able to independently adjust audio volume for each speaker. For instance, to mimic the siren, speaker **195** is turned to 80% volume; speaker **197** is adjusted to 40% volume; speaker **194** is set to 16% of volume; and speaker **196** is lowered to 8%.

[0063] Alternatively, when additional metadata becomes available, such data can also facilitate what is happening in a particular “place” which is associated with driver’s visual

field. For example, the visual field can be straight ahead, 30 degrees to the left, 60 degrees to the right, or coming from behind of the driver. The “spatially located” metadata, in one example, can be derived either from audio (multiple microphones) or synthesized by other sensors within the car. In another example, if one of the doors of a car is ajar or tapped, IAS or system can synthesize the “direction” where the offending door is. Once this positional metadata is known, the metadata can enhance user experiences.

[0064] A benefit of using IAS with spatially correlated audio and visual warnings is that IAS facilitates redirection driver or operator’s attention to the event direction when a vehicle emits a warning tone about something that is happening or is about to happen in a particular place (either inside or outside of the car). Since a vehicle has at least four independent speakers located throughout the cabin, these speakers can be driven with audio signals that “place” the warning tone anywhere in a 360-degree field within the car.

[0065] One advantage for spatially correlated audio warnings is that, by placing the audio warning in space, more information is conveyed. For example, for issuing a warning indicating which seat doesn’t have the seat belt fastened or which door has been tapped, IAS uses spatially correlated audio effect to sound a directional warning which identifies the origination of the event.

[0066] FIG. 1E illustrates block diagrams illustrating an ISA process of warning using positional metadata providing visual notification using a virtuous cycle in accordance with one embodiment of the present invention. Diagram **135** illustrates a steering wheel and a dashboard mounted inside of a vehicle. The dashboard, in one embodiment, includes a color-changing illuminable bar **136** amounted along the dashboard facing the operator. Color-changing illuminable bar **136** in diagram **135** emits a normal color light such as green light to indicate that the driver has scanned forward appropriately. Color-changing illuminable bar **136** emits a warning light **138** such as red light in diagram **137** to indicate that the driver has turned right but failed to check to the right before turning.

[0067] Diagram **250** including a left rear view mirror indicator, right rear view mirror indicator, and center rear view mirror indicator. All indicators in diagram **250** illuminate norm light indicating that the driver has been checking the mirrors at appropriate intervals. Most of the indicators in diagram **252** illuminate norm light except the right rear view mirror **256** which indicates in warning color such as red color which indicates that the driver has failed to check the right rear view mirror at appropriate intervals.

[0068] In one embodiment, IAS or system employs a multi-colored light bar across the front dash of a car. Normally the entire multi-colored light bar is green which indicates that in the full range of forward direction. The green bar indicates that the user or driver has focused his/her gaze and attention appropriately within a predefined time interval (i.e., last N seconds). As illustrated in diagram **137**, when a driver has failed to look at the right direction before making a right turn, the right side of the light bar can momentarily emit red instead of green. IAS provides an effect that a driver sees green under normal conditions). When the driver fails to check at the right direction before turning or fails to look forward for too long (i.e., texting), the driver will see a momentary red light bar. IAS using positional metadata to subtly “train” or “warn” the operator when he/she is failing to be properly attentive to the road.

[0069] Alternatively, IAS and/or system, in one embodiment, employs visual manifestation of positional metadata to provide a green/red light bar or rim of one or more rear-view mirrors. If the driver is checking them frequently enough, the bar would be green. If, however, the driver fails to check the mirrors enough times within a predefined interval, the bar can issue a notice by illuminating red light occasionally. The application of positional metadata to visually inform the driver can similarly be applied to the interior (non-road) components or vehicle. For example, a multi-colored light border can be installed around or adjacent to a navigation/infotainment system. In normal condition, the multi-colored light illuminates green light indicating current status. When the driver gazes or lingers at the navigation/infotainment system too long, the multi-colored light around the border of the navigation/infotainment system begins turning red to notifying the driver who should pay attention to the road. Note that the metadata object model of both interior and exterior context allows a sophisticated policy to be created on top of the object model. For instance, the policy may allow a driver to look at the navigation system for 3 seconds under some circumstances, but not under other circumstances such as when the vehicle is approaching a traffic light which is turning from yellow to red.

[0070] Automatic Adjustment of in-Car Entertainment Volume

[0071] IAS, in one embodiment, includes an auto adjustment capable of automatic adjusting in-car entertainment volume upon detecting an event. The volume adjustment component which uses one or more cameras and/or one or more microphones is able to extract real-time metadata regarding what is happening inside and outside of the vehicle. The in-car entertainment system volume can then be modulated or muted based on the real-time situation. For example, detection of an emergency vehicle siren (audio) or flashing lights (video) can trigger an automatic muting of the sound system which enables the driver and passengers to focus their attention appropriately. Other external situations, such as heavy rain and impending quick slowdown of traffic, could also be automatically detected (via metadata extraction from video and audio), IAS alerts the driver to focus exclusively on driving tasks by muting the entertainment system.

[0072] The following sample code logic illustrates one embodiment of how the real-time contextual object model that can be used to affect user-experience “policy” for a vehicle.

```
//
// hard rain -> max volume at 50%
//
if (CarObjectModel.exterior.raining.level > 0.8)
    setMaximumVolume(maxVolume * 0.5);
//
// mute audio if emergency vehicle is detected either visually or audibly
//
if (CarObjectModel.exterior.emergencyVehiclePresent)
    setMaximumVolume(0.0);
```

[0073] Audio Detection of Nail or Screw in a Tire

[0074] IAS, in one embodiment, includes an audio detector which includes multiple audio sensors capable of detecting audio sound relating to nail(s) or screw(s) in a tire. The system of audio detector, for example, obtains an audio data stream from a microphone placed near a tire wherein the

microphone(s) can monitor or detect the presence of a nail or screw embedded in the tire. Using frequency analysis of the audio and combining that with the current speed of the vehicle (via connection to OBD (On-board diagnostics) data bus), the presence of a nail or screw or other embedded object can be detected. Upon detecting the nail or screw, the system warns the driver in advance of total failure of the tire.

[0075] The audio detector, in one aspect, is able to detect rotation correlated audio anomalies. To detect anomalies, an audio metadata extracting software element takes digital sound samples and converts them via a FFT (fast Fourier Transform) into a series of coefficients representing the frequency components of the incoming audio signal. High frequency components that are modulated at the expected rate that would be found based on the circumference of the tire can be identified via either machine learning models or traditional pattern recognition software techniques. For example, if a tire has a circumference of C inches and the car is travelling at speed V (in miles per hour), the tire will rotate N times per second, where $N = (V * 5280 / 3600) / (C / 12)$. Thus, a high frequency audio anomaly would be expected to appear N times per second, and importantly, as the car changes velocity, the audio anomaly would shift its frequency of occurrence (N) to match. In order to identify a speed correlated audio anomaly, the audio data stream can be divided up into blocks of data whose size is calculated based on the expected length of time a single rotation of the tire will take. By calculating the degree of correlation between these data blocks using techniques such as the Pearson product-moment coefficient, the presence of a time correlated audio anomaly can be detected without concern over identifying the exact location within the sample that audio anomaly appears. Note that in addition to this technique, a machine learning model that has been trained to take audio and current velocity input and classify the audio as containing frequency dependent anomalies is also possible.

[0076] It should be noted that when a nail is embedded in a tire, it can result in a very slow leak. When that happens, external facing audio sensors or microphones placed in the wheel well of the vehicle can act similar as exterior facing cameras which are used to infer current context of the vehicle. In this example, the microphone sensors feed into a machine learning model that can classify audio signals as either being characteristic of having a nail in the tire or not. The classified metadata become available to the system. In one embodiment, the metadata can be used to inform the driver or operator about the nail and let the driver know that which tire has an embedded nail.

[0077] In another embodiment, the audio detector is able to predict how long and/or how many miles the vehicle can move before the tire becomes flat. The predication of failure will involve detected data, stored data, as well as models via a virtuous cycle.

[0078] An advantage of using the audio detector is to warn the driver one of the tires is about to be flat in a near future.

[0079] Predictive Part Failure

[0080] IAS includes a predictive failure component capable of predicting imminent device failure based on collected data as well as historical (stored) data. In one example, the predictive failure component uses multiple sensors such as microphones placed strategically on various components within a machine (such as a vehicle) to capture a data stream for the purpose of training machine learning

models. The captured audio data is frequency analyzed and collected in the aggregate along with other available sensor data. In addition, data collected based on scheduled and unscheduled maintenance (replacement of failed parts or components) is used to categorize the audio and sensor data. It should be noted that the data collected should enable the training of models that can identify component failure in advance of symptomatic failure.

[0081] In one aspect, the sound profile of a machine during operation can be used to determine the "health" of the system compared to a reference "known good" system. Any number of potential failures can be diagnosed via audible sound profiles much sooner than if one waits until the part or system fails. For example, when a belt break and thus fails to turn the alternator or water pump, the internal "sound" of the system will change. When brake pads wear down to the embedded metal used to detect wear, the "sound" of the system will change. All of these sounds are analyzed in the same or similar way that interior cameras are analyzed. The analysis via collected data as well as historical data can yield a real-time object model of how various components of the car are doing. The system can use this object model to craft user-experience that helps the driver to deal with impending failure of the vehicle.

[0082] Policy Fencing Based on Combined Sensor Data

[0083] IAS includes a policy fencing component capable of limiting certain activities upon detecting a set of pre-defined events. IAS or system using cameras, microphones, and other sensor inputs extracts real-time metadata regarding the current situation of an operator of a device or car. The system has a method of describing policy rules based on the input data and use those rules to identify situations where user or system constraints should be applied.

[0084] IAS or system collects data from the outside via direct sensor inputs (such as GPS, accelerometers, CAN bus data input) and also derives metadata from input sensor data such as video cameras and microphones. The extracted metadata includes various events, such as, but not limited to, "user is currently travelling in a car," "user is currently walking on a street with traffic going by at 35 mph," "audio input indicates an emergency vehicle is operating nearby," or "user is currently undergoing rapid acceleration."

[0085] The policy fencing component provides policy constraints in accordance with collected data and a set of predefined policies. The policies, in one aspect, are both declarative and procedural, and allow for the real-time derivation of "current system constraints" that should be applied to both the system and operators. For example, a policy can be defined that "operator cannot play current video game because current velocity exceeds 12 miles per hour", but that policy can be refined such that, if the system can tell via video camera input that the operator is a passenger of a vehicle and not the driver, the constraint does not apply. Another example could be a policy for the player of an augmented reality mobile application game. The system can decide that when the user is near traffic, or in the crosswalk of a street, game play would not be allowed. Another example could be simple geo-fencing, where a bounded set by GPS location which defines an area where no game play is allowed.

[0086] It should be noted that combination of all sensor input and extracted metadata can be used to define a set of policies that depend on more than a single input constraint. Note that historical metadata, such as how long a given

activity has been conducted, will be part of the data stream. For example, the constraints can be that “user can’t do X once they’ve exceeded Y hours of play, unless they are within location Z, in which case they are allowed to do X.”

[0087] Required Maintenance Detection

[0088] IAS includes a required maintenance finder capable of identifying potential objects needed repair or maintenance based on collected data as well as historical (stored) data. For example, IAS is able to extract metadata from a fleet of moving vehicles to identify and schedule for repair roads, building, or other infrastructure in need of repair. For example, when a car goes over a pothole, the car could collect data about the length and severity of the hole. A fleet of such cars, sending this information into the cloud, will allow municipalities or other interested parties to identify needed repairs. Emergency situations, such as when a load of building materials has fallen off a truck onto a busy highway, can also be identified by the pattern of swerving of the vehicles, as well as by extracted metadata from an outward facing video camera as described in many scenarios above.

[0089] In addition to using all of the generated real-time metadata in order to deliver a better user-experience within the car, the metadata is also transmitted to a cloud service for aggregation and analysis. The aggregated data can carry additional meaning that cannot be found from a single vehicle.

[0090] FIG. 1F is a logic block diagram illustrating a pipeline process 150 of outward facing camera capable of identifying and classifying detected object(s) using a virtuous cycle in accordance with one embodiment of the present invention. Outward facing camera 151 collects images and the images are stored in a queue 152. After scaling the images by image scaling component 153, the scaled image is forwarded to object detection 154. Object detection 154 generates a collection of objection information which is forwarded to queue 155. The object information which includes bounding-box, object category, object orientation, and object distance is forwarded to component 156 and router 157. Upon categorizing the object information at block 156, the categorized data is forwarded to map 158. After recognizing the object based on map 158, the recognizer is forwarded to router 157. After routing information at router 157, the output images are forwarded to block 159 which uses classifier 130-131 to classify the images and/or objects.

[0091] Pipeline process 150 illustrates a logic processing flow which is instantiated for the purpose of processing incoming data, extracting metadata on a frame by frame or data packet basis, and forwarding both frames and metadata packets forward through the pipeline. Each stage of the pipeline can contain software elements that perform operations upon the current audio or video or sensor data frame. The elements in the pipeline can be inserted or removed while the pipeline is running, which allows for an adaptive pipeline that can perform different operations depending on the applications. The pipeline process is configured to adapt various system constraints that can be situationally present. Additionally, elements in the pipeline can have their internal settings updated in real-time, providing the ability to “turn off,” “turn on” elements, or to adjust their configuration settings on the fly. For example, a given element in a pipeline can have a setting such that it operates on every 10th frame in the data stream, and such setting can be

changed in real-time to adjust to a different frame rate. Elements in the data stream can also emit newly constructed data packets into the stream. In one embodiment, the pipeline process 150 allows extraction of higher level meaning from the stream to be forwarded onto software components that deal with higher level events. For example, a pipeline element might be able to recognize a “stop light” in a frame of video, and make a determination that in that frame it is current state is “yellow”. Such an element would construct a metadata packet containing that information, and subsequent downstream software components would receive that packet and be able to act upon it.

[0092] Pipeline process 150 includes a metadata packet schema which includes name/value pairs with arbitrary nesting and basic primitive data types such as arrays and structures that is used to create a self-describing and both machine and human readable form of the extracted real-time metadata flowing through the system. Such a generalized schema allows multiple software components to agree on how to describe the high level events that are being captured and analyzed and acted upon by the system. For example, a schema is constructed to describe the individual locations within a video frame of a person’s eyes, nose, mouth, chin line, etc. Such a data structure allows a downstream software component to infer even higher level events, such as “this person is looking up at 34 degrees above the horizon” or “this person is looking left 18 degrees left of center.” The process can subsequently construct additional metadata packets and insert them into the stream, resulting in higher level semantic metadata that the system is able to act upon.

[0093] FIG. 1G is a logic block diagram illustrating a pipeline process 160 capable of identifying and classifying face detection, head and gaze orientation, and mouth features using a virtuous cycle in accordance with one embodiment of the present invention. Inward facing camera 161 collects images and the images are stored in a queue 162. After scaling the images by image scaling component 163, the scaled image is forwarded to face and head detection 164. The output of detection 164 is forwarded to image transform (“IT”) components 165-166. After transformation, the transformed image is forwarded to blocks 169-170. After facial feature extraction in block 169, the feature map is forwarded to block 167 for pose normalization. Block 168 receives face images from IT component 165 and transformed images from block 167, the normalized face image is forwarded to block 172. Upon processing normalized face with embedding network at block 172, a face ID is identified.

[0094] Block 170 extracts mouth feature and generates mouth feature(s) of driver. Block 171 processes head and gaze based on output of IT component 166 which receives information with both scaled and unscaled images. In one example, block 171 is capable of generating various features, such as gaze, head, number of eyes, glasses, and the like.

[0095] Pipeline process 160, in one example, includes a process of automatic entertainment sound system volume adjustment. For example, a data pipeline of audio, video, and sensor data flowing through a series of software elements, one type of real-time extracted metadata is the presence of emergency vehicles in the vicinity. An audio metadata extracting software element takes digital sound samples and converts them via a FFT (fast Fourier Transform) into a series of coefficients representing the frequency

components of the incoming audio signal. Such data can be analyzed for time-series fingerprints that are typical of emergency vehicle sirens, which consist of frequency ramps and alternating a pair of frequencies. This data can also be provided as input into a machine learning classification model that can identify sound patterns such as sirens. Note also that the “doppler shift” effect is factored into these recognizer software elements, because moving vehicles relative to a listener will have a frequency shift that is a function of the speed and direction of relative motion. In addition to audio siren detection, video frames that are flowing through the pipeline can detect the presence of flashing lights that are typically used by emergency vehicles. When the presence of an active emergency vehicle is detected in the vicinity, the sound system is muted, which will allow the driver to become aware of the emergency vehicle.

[0096] In one aspect, pipeline process 160 further includes a process of attaching various microphones to various internal components. For example, data pipeline of audio, video, and sensor data flowing through a series of software elements, a multitude of audio sources can be obtained by placing microphones at various points in the internal mechanisms of a machine. In an automobile, microphones could be placed in each wheel well, motor, and/or various other points within the engine and body frame. The audio from each of these microphones can pick up vibrational pattern of attached devices during operation. An audio metadata extracting software element takes digital sound samples and converts them via a FFT (fast Fourier Transform) into a series of coefficients representing the frequency components of incoming audio signal. The FFT encoded signal pattern can be uploaded to a cloud-based data repository, where, in aggregate with data sets collected from other machines of like type, a machine learning neural-net can be trained to recognize normal and abnormal patterns. In addition, when the machine undergoes either scheduled or unscheduled maintenance, those maintenance events will serve as classification data to the encoded audio streams. In one aspect, the process will enable the machine learning component to “predict” imminent or incipient part failure. For example, if a car manufacturer were to apply this technique to hundreds of thousands of vehicles, patterns that precede part failure could be predicted in advance. Imagine that a rotating part that is subject to failure tends to create a certain vibrational pattern well in advance of catastrophic failure. That vibrational pattern would be picked up by an attached microphone, the digitized audio signal would be fed into a machine learning model capable of detecting such patterns. This model, when applied to a machine that has not yet suffered catastrophic failure, would be able to issue a warning indicating that the part should be examined and possibly replaced in advance of failure.

[0097] FIG. 1H is a logic block diagram 175 illustrating a process of classifying detected object(s) using a virtuous cycle in accordance with one embodiment of the present invention. Block 176 is a software element used to classify a pedestrian based on collected external images captured by outward facing cameras. Based on collected data and historical data, pedestrian may be identified. Block 177 is a software element used to classify a vehicle based on collected external images captured by outward facing cameras. Based on collected data and historical data, vehicle information can be identified. The exemplary classification information includes model of the vehicle, license plate, state of

vehicle registration, and the like. In addition, information such as turn-signals, brake lights, and headlights can also be classified via facilitation of virtuous cycle. Block 178 is a software element used to classify traffic signals or conditions according to collected external images captured by outward facing cameras. For example, according to collected data as well as historical data, the traffic signal can be classified. The exemplary classification includes sign, speed limit, stop sign, and the like.

[0098] FIG. 2A is a block diagram 200 illustrating a virtuous cycle capable of detecting or monitoring IA in accordance with one embodiment of the present invention. Diagram 200, which is similar to diagram 100 shown in FIG. 1A, includes a containerized sensor network 206, real-world scale data 202, and continuous machine learning 204. In one embodiment, continuous machine learning 204 pushes real-time models to containerized sensor network 206 as indicated by numeral 210. Containerized sensor network 206 continuously feeds captured data or images to real-world scale data 202 with uploading in real-time or in a batched format. Real-world scale data 202 provides labeled data to continuous machine learning 204 for constant model training as indicated by numeral 212. It should be noted that the underlying concept of the exemplary embodiment(s) of the present invention would not change if one or more blocks (or elements) were added to or removed from FIG. 2A.

[0099] The virtuous cycle illustrated in diagram 200, in one embodiment, is configured to implement IAS wherein containerized sensor network 206 is similar to vehicle 102 as shown in FIG. 1A and real-world scale data 202 is similar to CBN 104 shown in FIG. 1A. Also, continuous machine learning 204 is similar to MCL 106 shown in FIG. 1A. In one aspect, containerized sensor network 206 such as an automobile or car contains a containerized sensing device capable of collecting surrounding information or images using onboard sensors or sensor network when the car is in motion. Based on the IA model, selective recording of the collected surrounding information is selectively recorded to a local storage or memory.

[0100] Real-world scale data 202, such as cloud or CBN, which is wirelessly coupled to the containerized sensing device, is able to correlate with cloud data and recently obtained IA data for producing labeled data. For example, real-world scale data 202 generates IA labeled data based on historical IA cloud data and the surrounding information sent from the containerized sensing device.

[0101] Continuous machine learning 204, such as MLC or cloud, is configured to train and improve IA model based on the labeled data from real-world scale data 202. With continuous gathering data and training IA model(s), the IAS will be able to learn, obtain, and/or collect all available IAS for the population samples.

[0102] In one embodiment, a virtuous cycle includes partitionable Machine Learning networks, training partitioned networks, partitioning a network using sub-modules, and composing partitioned networks. For example, a virtuous cycle involves data gathering from a device, creating intelligent behaviors from the data, and deploying the intelligence. In one example, partition idea includes knowing the age of a driver which could place or partition “dangerous driving” into multiple models and selectively deployed by an “age detector.” An advantage of using such partitioned models is that models should be able to perform a better job of recognition with the same resources because the domain

of discourse is now smaller. Note that, even if some behaviors overlap by age, the partitioned models can have common recognition components.

[0103] It should be noted that more context information collected, a better job of recognition can be generated. For example, “dangerous driving” can be further partitioned by weather condition, time of day, traffic conditions, et cetera. In the “dangerous driving” scenario, categories of dangerous driving can be partitioned into “inattention”, “aggressive driving”, “following too closely”, “swerving”, “driving too slowly”, “frequent breaking”, deceleration, ABS event, et cetera.

[0104] For example, by resisting a steering behavior that is erratic, the car gives the driver direct feedback on their behavior—if the resistance is modest enough then if the steering behavior is intentional (such as trying to avoid running over a small animal) then the driver is still able to perform their irregular action. However, if the driver is texting or inebriated then the correction may alert them to their behavior and get their attention. Similarly, someone engaged in “road rage” who is driving too close to another car may feel resistance on the gas pedal. A benefit of using IAS is to identify consequences of a driver’s “dangerous behavior” as opposed to recognizing the causes (texting, etc.). The Machine Intelligence should recognize the causes as part of the analysis for offering corrective action.

[0105] In one aspect, a model such as IA model includes some individual blocks that are trained in isolation to the larger problem (e.g. weather detection, traffic detection, road type, etc.). Combining the blocks can produce a larger model. Note that the sample data may include behaviors that are clearly bad (ABS event, rapid deceleration, midline crossing, being too close to the car in front, etc.). In one embodiment, one or more sub-modules are built. The models include weather condition detection and traffic detection for additional modules intelligence, such as “correction vectors” for “dangerous driving.”

[0106] An advantage of using a virtuous cycle is that it can learn and detect object such as IA in the real world.

[0107] FIG. 2B is a block diagram 230 illustrating an alternative exemplary virtuous cycle capable of detecting IA in accordance with one embodiment of the present invention. Diagram 230 includes external data source 234, sensors 238, crowdsourcing 233, and intelligent model 239. In one aspect, components/activities above dotted line 231 are operated in cloud 232, also known as in-cloud component. Components/activities below dotted line 231 are operated in car 236, also known as in-device or in-car component. It should be noted that the underlying concept of the exemplary embodiment(s) of the present invention would not change if one or more blocks (or elements) were added to or removed from FIG. 2B.

[0108] In one aspect, in-cloud components and in-device components coordinate to perform desirable user specific tasks. While in-cloud component leverages massive scale to process incoming device information, cloud applications leverage crowd sourced data to produce applications. External data sources can be used to contextualize the applications to facilitate intellectual crowdsourcing. For example, in-car (or in-phone or in-device) portion of the virtuous cycle pushes intelligent data gathering to the edge application. In one example, edge applications can perform intelligent data gathering as well as intelligent in-car processing.

It should be noted that the amount of data gathering may rely on sensor data as well as intelligent models which can be loaded to the edge.

[0109] FIG. 3 is a block diagram 300 illustrating a cloud based network using crowdsourcing approach to improve IA model(s) in accordance with one embodiment of the present invention. Diagram 300 includes population of vehicles 302, sample population 304, models deployment 306, correlation component 308, and cloud application 312. It should be noted that the underlying concept of the exemplary embodiment(s) of the present invention would not change if one or more blocks (or samples) were added to or removed from FIG. 3.

[0110] Crowdsourcing is a process of using various sourcing or specific models generated or contributed from other cloud or Internet users for achieving needed services. For example, crowdsourcing relies on the availability of a large population of vehicles, phones, or other devices to source data 302. For example, a subset of available devices such as sample 304 is chosen by some criterion such as location to perform data gathering tasks. To gather data more efficiently, intelligent models are deployed to a limited number of vehicles 306 for reducing the need of large uploading and processing a great deal of data in the cloud. It should be noted that the chosen devices such as cars 306 monitor the environment with the intelligent model and create succinct data about what has been observed. The data generated by the intelligent models is uploaded to the correlated data store as indicated by numeral 308. It should be noted that the uploading can be performed in real-time for certain information or at a later time for other types of information depending on the need as well as condition of network traffic.

[0111] Correlated component 308 includes correlated data storage capable of providing a mechanism for storing and querying uploaded data. Cloud applications 312, in one embodiment, leverage the correlated data to produce new intelligent models, create crowd sourced applications, and other types of analysis.

[0112] FIG. 4 is a block diagram 400 illustrating an IAS using the virtuous cycle in accordance with one embodiment of the present invention. Diagram 400 includes a correlated data store 402, machine learning framework 404, and sensor network 406. Correlated data store 402, machine learning framework 404, and sensor network 406 are coupled by connections 410-416 to form a virtuous cycle as indicated by numeral 420. It should be noted that the underlying concept of the exemplary embodiment(s) of the present invention would not change if one or more blocks (circuit or elements) were added to or removed from FIG. 4.

[0113] In one embodiment, correlated data store 402 manages real-time streams of data in such a way that correlations between the data are preserved. Sensor network 406 represents the collection of vehicles, phones, stationary sensors, and other devices, and is capable of uploading real-time events into correlated data store 402 via a wireless communication network 412 in real-time or in a batched format. In one aspect, stationary sensors includes, but not limited to, municipal cameras, webcams in offices and buildings, parking lot cameras, security cameras, and traffic cams capable of collecting real-time images.

[0114] The stationary cameras such as municipal cameras and webcams in offices are usually configured to point to streets, buildings, parking lots wherein the images captured

by such stationary cameras can be used for accurate labeling. To fuse between motion images captured by vehicles and still images captured by stationary cameras can track object (s) such as car(s) more accurately. Combining or fusing stationary sensors and vehicle sensors can provide both labeling data and historical stationary sampling data also known as stationary “fabric”. It should be noted that during the crowdsourcing applications, fusing stationary data (e.g. stationary cameras can collect vehicle speed and position) with real-time moving images can improve ML process.

[0115] Machine Learning (“ML”) framework 404 manages sensor network 406 and provides mechanisms for analysis and training of ML models. ML framework 404 draws data from correlated data store 402 via a communication network 410 for the purpose of training modes and/or labeled data analysis. ML framework 404 can deploy data gathering modules to gather specific data as well as deploy ML models based on the previously gathered data. The data upload, training, and model deployment cycle can be continuous to enable continuous improvement of models.

[0116] FIG. 5 is a block diagram 500 illustrating an exemplary process of correlating IA data in accordance with one embodiment of the present invention. Diagram 500 includes source input 504, real-time data management 508, history store 510, and crowd sourced applications 512-516. In one example, source input 504 includes cars, phones, tablets, watches, computers, and the like capable of collecting massive amount of data or images which will be passed onto real-time data management 508 as indicated by numeral 506. It should be noted that the underlying concept of the exemplary embodiment(s) of the present invention would not change if one or more blocks (or elements) were added to or removed from FIG. 5.

[0117] In one aspect, a correlated system includes a real-time portion and a batch/historical portion. The real-time part aims to leverage new data in near or approximately real-time. Real-time component or management 508 is configured to manage a massive amount of influx data 506 coming from cars, phones, and other devices 504. In one aspect, after ingesting data in real-time, real-time data management 508 transmits processed data in bulk to the batch/historical store 510 as well as routes the data to crowd sourced applications 512-516 in real-time.

[0118] Crowd sourced applications 512-516, in one embodiment, leverage real-time events to track, analyze, and store information that can be offered to user, clients, and/or subscribers. Batch-Historical side of correlated data store 510 maintains a historical record of potentially all events consumed by the real-time framework. In one example, historical data can be gathered from the real-time stream and it can be stored in a history store 510 that provides high performance, low cost, and durable storage. In one aspect, real-time data management 508 and history store 510 coupled by a connection 502 are configured to perform IA data correlation as indicated by dotted line.

[0119] FIG. 6 is a block diagram 600 illustrating an exemplary process of real-time data for IAS in accordance with one embodiment of the present invention. Diagram 600 includes data input 602, gateway 606, normalizer 608, queue 610, dispatcher 616, storage conversion 620, and historical data storage 624. The process of real-time data management further includes a component 614 for publish and subscribe. It should be noted that the underlying concept of the exemplary embodiment(s) of the present invention would

not change if one or more blocks (circuit or elements) were added to or removed from FIG. 6.

[0120] The real-time data management, in one embodiment, is able to handle a large numbers (i.e., 10’s of millions) of report events to the cloud as indicated by numeral 604. API (application program interface) gateway 606 can handle multiple functions such as client authentication and load balancing of events pushed into the cloud. The real-time data management can leverage standard HTTP protocols. The events are routed to stateless servers for performing data scrubbing and normalization as indicated by numeral 608. The events from multiple sources 602 are aggregated together into a scalable/durable/consistent queue as indicated by numeral 610. An event dispatcher 616 provides a publish/subscribe model for crowd source applications 618 which enables each application to look at a small subset of the event types. The heterogeneous event stream, for example, is captured and converted to files for long-term storage as indicated by numeral 620. Long-term storage 624 provides a scalable and durable repository for historical data.

[0121] FIG. 7 is a block diagram 700 illustrating a crowd sourced application model for IA in accordance with one embodiment of the present invention. Diagram 700 includes a gateway 702, event handler 704, state cache 706, state store 708, client request handler 710, gateway 712, and source input 714. In one example, gateway 702 receives an event stream from an event dispatcher and API gateway 712 receives information/data from input source 714. It should be noted that the underlying concept of the exemplary embodiment(s) of the present invention would not change if one or more blocks (or elements) were added to or removed from FIG. 7.

[0122] The crowd sourced application model, in one embodiment, facilitates events to be routed to a crowd source application from a real-time data manager. In one example, the events enter gateway 702 using a simple push call. Note that multiple events are handled by one or more servers. The events, in one aspect, are converted into inserts or modifications to a common state store. State store 708 is able to hold data from multiple applications and is scalable and durable. For example, State store 708, besides historical data, is configured to store present data, information about “future data”, and/or data that can be shared across applications such as predictive AI (artificial intelligence).

[0123] State cache 706, in one example, is used to provide fast access to commonly requested data stored in state store 708. Note that application can be used by clients. API gateway 712 provides authentication and load balancing. Client request handler 710 leverages state store 708 for providing client data.

[0124] In an exemplary embodiment, an onboard IA model is able to handle real-time IA detection based on triggering events. For example, after ML models or IA models for IA detection have been deployed to all or most of the vehicles, the deployed ML models will report to collected data indicating IAS for facilitating issuance of real-time warning for dangerous event(s). The information or data relating to the real-time dangerous event(s) or IAS is stored in state store 708. Vehicles 714 looking for IA detection can, for example, access the IAS using gateway 712.

[0125] FIG. 8 is a block diagram 800 illustrating a method of storing IA related data using a geo-spatial objective storage in accordance with one embodiment of the present

invention. Diagram **800** includes gateway **802**, initial object **804**, put call **806**, find call **808**, get call **810**, SQL (Structured Query Language) **812**, non-SQL **814**, and geo-spatial object storage **820**. It should be noted that the underlying concept of the exemplary embodiment(s) of the present invention would not change if one or more blocks (circuit or elements) were added to or removed from FIG. **8**.

[**0126**] Geo-spatial object storage **820**, in one aspect, stores or holds objects which may include time period, spatial extent, ancillary information, and optional linked file. In one embodiment, geo-spatial object storage **820** includes UUID (universally unique identifier) **822**, version **824**, start and end time **826**, bounding **828**, properties **830**, data **832**, and file-path **834**. For example, while UUID **822** identifies an object, all objects have version(s) **824** that allow schema to change in the future. Start and end time **826** indicates an optional time period with a start time and an end time. An optional bounding geometry **828** is used to specify spatial extent of an object. An optional set of properties **830** is used to specify name-value pairs. Data **832** can be binary data. An optional file path **834** may be used to associate with the object of a file containing relevant information such as MPEG (Moving Picture Experts Group) stream.

[**0127**] In one embodiment, API gateway **802** is used to provide access to the service. Before an object can be added to the store, the object is assigned an UUID which is provided by the initial object call. Once UUID is established for a new object, the put call **804** stores the object state. The state is stored durably in Non-SQL store **814** along with UUID. A portion of UUID is used as hash partition for scale-out. The indexable properties includes version, time duration, bounding, and properties which are inserted in a scalable SQL store **812** for indexing. The Non-SQL store **814** is used to contain the full object state. Non-SQL store **814** is scaled-out using UUID as, for example, a partition key.

[**0128**] SQL store **812** is used to create index tables that can be used to perform queries. SQL store **812** may include three tables **816** containing information, bounding, and properties. For example, information holds a primary key, objects void, creation timestamp, state of object and object properties “version” and “time duration.” Bounding holds the bounding geometry from the object and the id of the associated information table entry. Properties hold property name/value pairs from the object stored as one name/value pair per row along with ID of associated info table entry.

[**0129**] Find call **808**, in one embodiment, accepts a query and returns a result set, and issues a SQL query to SQL store **812** and returns a result set containing UUID that matches the query.

[**0130**] FIG. **9** is a block diagram **900** illustrating an exemplary approach of analysis engine analyzing IA data in accordance with one embodiment of the present invention. Diagram **900** includes history store **902**, analysis engine **904**, and geo-spatial object store **906**. It should be noted that the underlying concept of the exemplary embodiment(s) of the present invention would not change if one or more blocks (circuit or elements) were added to or removed from FIG. **9**.

[**0131**] In one aspect, diagram **900** illustrates analysis engine **904** containing ML training component capable of analyzing labeled data based on real-time captured IA data and historical data. The data transformation engine, in one example, interacts with Geo-spatial object store **906** to

locate relevant data and with history store to process the data. Optimally, the transformed data may be stored.

[**0132**] It should be noted that virtuous cycle employing ML training component to provide continuous model training using real-time data as well as historical samples, and deliver IA detection model for one or more subscribers. A feature of virtuous cycle is able to continuous training a model and able to provide a real-time or near real-time result. It should be noted that the virtuous cycle is applicable to various other fields, such as, but not limited to, business intelligence, law enforcement, medical services, military applications, and the like.

[**0133**] FIG. **10** is a block diagram **1000** illustrating an exemplary containerized sensor network used for sensing IAS related information in accordance with one embodiment of the present invention. Diagram **1000** includes a sensor bus **1002**, streaming pipeline **1004**, and application layer **1006** wherein sensor bus **1002** is able to receive low-bandwidth sources and high-bandwidth sources. Streaming pipeline **1004**, in one embodiment, includes ML capable of generating unique model such as model **1008**. It should be noted that the underlying concept of the exemplary embodiment(s) of the present invention would not change if one or more blocks (circuit or elements) were added to or removed from FIG. **10**.

[**0134**] FIG. **11** is a block diagram **1100** illustrating a processing device or computer system which can be installed in a vehicle to support onboard cameras, CAN (Controller Area Network) bus, Inertial Measurement Units, Lidar, et cetera for facilitating virtuous cycle in accordance with one embodiment of the present invention. Computer system or IAS **1100** can include a processing unit **1101**, an interface bus **1112**, and an input/output (“IO”) unit **1120**. Processing unit **1101** includes a processor **1102**, a main memory **1104**, a system bus **1111**, a static memory device **1106**, a bus control unit **1105**, I/O element **1130**, and IAS element **1185**. It should be noted that the underlying concept of the exemplary embodiment(s) of the present invention would not change if one or more blocks (circuit or elements) were added to or removed from FIG. **11**.

[**0135**] Bus **1111** is used to transmit information between various components and processor **1102** for data processing. Processor **1102** may be any of a wide variety of general-purpose processors, embedded processors, or microprocessors such as ARM® embedded processors, Intel® Core™ Duo, Core™ Quad, Xeon®, Pentium™ microprocessor, Motorola™ 68040, AMD® family processors, or Power PC™ microprocessor.

[**0136**] Main memory **1104**, which may include multiple levels of cache memories, stores frequently used data and instructions. Main memory **1104** may be RAM (random access memory), MRAM (magnetic RAM), or flash memory. Static memory **1106** may be a ROM (read-only memory), which is coupled to bus **1111**, for storing static information and/or instructions. Bus control unit **1105** is coupled to buses **1111-1112** and controls which component, such as main memory **1104** or processor **1102**, can use the bus. Bus control unit **1105** manages the communications between bus **1111** and bus **1112**.

[**0137**] I/O unit **1120**, in one embodiment, includes a display **1121**, keyboard **1122**, cursor control device **1123**, and communication device **1125**. Display device **1121** may be a liquid crystal device, cathode ray tube (“CRT”), touch-screen display, or other suitable display device. Display

1121 projects or displays images of a graphical planning board. Keyboard **1122** may be a conventional alphanumeric input device for communicating information between computer system **1100** and computer operator(s). Another type of user input device is cursor control device **1123**, such as a conventional mouse, touch mouse, trackball, or other type of cursor for communicating information between system **1100** and user(s).

[0138] IA element **1185**, in one embodiment, is coupled to bus **1111**, and configured to interface with the virtuous cycle for facilitating IA detection(s). For example, if IAS **1100** is installed in a car, IA element **1185** is used to operate the IA model as well as interface with the cloud based network. If IAS **1100** is placed at the cloud based network, IA element **1185** can be configured to handle the correlating process for generating labeled data.

[0139] Communication device **1125** is coupled to bus **1111** for accessing information from remote computers or servers, such as server **104** or other computers, through wide-area network **102**. Communication device **1125** may include a modem or a network interface device, or other similar devices that facilitate communication between computer **1100** and the network. Computer system **1100** may be coupled to a number of servers via a network infrastructure such as the Internet.

[0140] The exemplary embodiment of the present invention includes various processing steps, which will be described below. The steps of the embodiment may be embodied in machine or computer executable instructions. The instructions can be used to cause a general purpose or special purpose system, which is programmed with the instructions, to perform the steps of the exemplary embodiment of the present invention. Alternatively, the steps of the exemplary embodiment of the present invention may be performed by specific hardware components that contain hard-wired logic for performing the steps, or by any combination of programmed computer components and custom hardware components.

[0141] FIG. **12** is a flowchart **1200** illustrating a process of IAS capable of assisting vehicle operation via metadata extraction processed by IA model in accordance with one embodiment of the present invention. At block **1202**, the process activating a set of outward facing cameras mounted on a vehicle for recording external surrounding images representing a geographic environment in which the vehicle operates. At block **1204**, at least one of a set of inward facing cameras mounted in the vehicle is selectively enabled for collecting interior images of the vehicle. At block **1206**, the process identifies a targeted direction focused by operator eyes in accordance with the interior images and stored data managed by the IA model. At block **1208**, an external target is determined in response to the external surrounding images and the targeted direction. At block **1210**, the process identifies an object that the operator eyes are looking based on the external target and information supplied by the IA model.

[0142] In one embodiment, the process is further capable of tracking surrounding environmental event in accordance with the external surrounding images and information provided by the IA model trained by a virtuous cycle. For example, upon identifying potential vehicle collision in response to the surrounding environmental event, a spatially correlated audio warning is issued or provided to warn the operator if the potential vehicle collision is determined.

Alternatively, automatic adjustment of interior entertainment volume can also be provided or operated when the spatially correlated audio warning is activated. The process is further able to adjust moving direction of the vehicle to avoid potential collision in accordance with a set of pre-defined collision avoidance policy facilitated by the IA model which is trained by a virtuous cycle via a cloud computing.

[0143] While particular embodiments of the present invention have been shown and described, it will be obvious to those of ordinary skills in the art that based upon the teachings herein, changes and modifications may be made without departing from this exemplary embodiment(s) of the present invention and its broader aspects. Therefore, the appended claims are intended to encompass within their scope all such changes and modifications as are within the true spirit and scope of this exemplary embodiment(s) of the present invention.

What is claimed is:

1. A method comprising:
 - receiving audio data from a microphone positioned to capture sounds of a tire on a vehicle moving on a road;
 - obtaining a speed of the vehicle that overlaps with a same timeframe of when the sounds of the tire were captured by the microphone in the received audio data;
 - determining that a foreign object is embedded in the tire based on a frequency analysis of the received audio data relative to the speed of the vehicle; and
 - outputting an alert to a driver of the vehicle indicating the foreign object is embedded in the tire.
2. The method of claim 1, wherein determining that the foreign object is embedded in the tire includes:
 - selecting a plurality of sound samples from the received audio data;
 - converting the plurality of sound samples into a series of coefficients representing frequency components of the received audio data;
 - determining an expected rate of rotation for the tire based on the speed of the vehicle;
 - in response to high frequency components of the series of coefficients being modulated with respect to the expected rate of rotation, determining that the foreign object is embedded in the tire; and
 - in response to a lack of high frequency components of the series of coefficients being modulated with respect to the expected rate of rotation, determining that no foreign object is embedded in the tire.
3. The method of claim 1, wherein determining that the foreign object is embedded in the tire includes:
 - determining a rotation rate of the tire based on the speed and a circumference of the tire;
 - identifying an audio anomaly in the received audio data occurring at the determined rotation rate; and
 - determining that the foreign object is embedded in the tire in response to identification of the audio anomaly.
4. The method of claim 1, wherein determining that the foreign object is embedded in the tire includes:
 - dividing the received audio data into data blocks whose size is calculated based on an expected length of time of a single rotation of the tire;
 - detecting an audio anomaly in the received audio data based on a degree of correlation between the data blocks; and

- determining that the foreign object is embedded in the tire based on the audio anomaly.
- 5.** The method of claim **1**, wherein determining that the foreign object is embedded in the tire includes:
- performing a Pearson product-moment coefficient technique on the received audio data to detect a time correlated audio anomaly indicative of the foreign object being embedded in the tire.
- 6.** The method of claim **1**, wherein determining that the foreign object is embedded in the tire includes:
- employing a trained machine learning model on the received audio data to classify the received audio data as containing a frequency dependent anomaly indicative of the foreign object being embedded in the tire.
- 7.** The method of claim **1**, wherein determining that the foreign object is embedded in the tire includes:
- determining that a nail or screw object is embedded in the tire based on the frequency analysis of the received audio data relative to the speed of the vehicle.
- 8.** A system comprising:
- a microphone configured to capture sounds of a tire on a vehicle moving on a road;
 - an output interface configured to present an alert to a driver of the vehicle;
 - a memory configured to store computer instructions;
 - at least one processor configured to execute the computer instructions to:
 - receive, via the microphone, audio data for a time period;
 - obtain a speed of the vehicle during the time period;
 - determine that a foreign object is embedded in the tire based on a frequency analysis of the audio data relative to the speed of the vehicle; and
 - output, via the output interface, an alert to a driver of the vehicle indicating the foreign object is embedded in the tire.
- 9.** The system of claim **8**, wherein the processor is configured to determine that the foreign object is embedded in the tire by executing further computer instructions to:
- generate a plurality of sound samples from the audio data;
 - convert the plurality of sound samples into a series of coefficients representing frequency components of the audio data;
 - determine an expected rate of rotation for the tire based on the speed of the vehicle;
 - in response to high frequency components of the series of coefficients being modulated with respect to the expected rate of rotation, determine that the foreign object is embedded in the tire; and
 - in response to a lack of high frequency components of the series of coefficients being modulated with respect to the expected rate of rotation, determine that no foreign object is embedded in the tire.
- 10.** The system of claim **8**, wherein the processor is configured to determine that the foreign object is embedded in the tire by executing further computer instructions to:
- determine a rotation rate of the tire based on the speed and a circumference of the tire;
 - identify an audio anomaly in the audio data occurring at the determined rotation rate; and
 - determine that the foreign object is embedded in the tire in response to identification of the audio anomaly.
- 11.** The system of claim **8**, wherein the processor is configured to determine that the foreign object is embedded in the tire by executing further computer instructions to:
- divide the audio data into data blocks whose size is calculated based on an expected length of time of a single rotation of the tire;
 - detect an audio anomaly in the audio data based on a degree of correlation between the data blocks; and
 - determine that the foreign object is embedded in the tire based on the audio anomaly.
- 12.** The system of claim **8**, wherein the processor is configured to determine that the foreign object is embedded in the tire by executing further computer instructions to:
- perform a Pearson product-moment coefficient technique on the audio data to detect a time correlated audio anomaly indicative of the foreign object being embedded in the tire.
- 13.** The system of claim **8**, wherein the processor is configured to determine that the foreign object is embedded in the tire by executing further computer instructions to:
- employ a trained machine learning model on the audio data to classify the audio data as containing a frequency dependent anomaly indicative of the foreign object being embedded in the tire.
- 14.** The system of claim **8**, wherein the processor is configured to determine that the foreign object is embedded in the tire by executing further computer instructions to:
- determine that a nail or screw object is embedded in the tire based on the frequency analysis of the audio data relative to the speed of the vehicle.
- 15.** A computing device comprising:
- a memory configured to store computer instructions; and
 - at least one processor configured to execute the computer instructions to:
 - obtain audio data from a microphone positioned to capture sounds of a tire on a vehicle moving on a road;
 - determine a speed of the vehicle that corresponds to when the microphone captured the sounds of the tire;
 - determine that an object is embedded in the tire based on a frequency analysis of the audio data relative to the speed of the vehicle; and
 - output, via an output interface, an alert to a driver of the vehicle indicating the object is embedded in the tire.
- 16.** The computing device of claim **15**, wherein the processor is configured to determine that the object is embedded in the tire by executing further computer instructions to:
- select a plurality of sound samples from the audio data;
 - convert the plurality of sound samples into a series of coefficients representing frequency components of the audio data;
 - determine an expected rate of rotation for the tire based on the current speed of the vehicle;
 - in response to high frequency components of the series of coefficients being modulated with respect to the expected rate of rotation, determine that the object is embedded in the tire; and
 - in response to a lack of high frequency components of the series of coefficients being modulated with respect to the expected rate of rotation, determine that no object is embedded in the tire.

17. The computing device of claim **15**, wherein the processor is configured to determine that the object is embedded in the tire by executing further computer instructions to:

- determine a rotation rate of the tire based on the speed and a circumference of the tire;
- identify an audio anomaly in the audio data occurring at the determined rotation rate; and
- determine that the object is embedded in the tire in response to identification of the audio anomaly.

18. The computing device of claim **15**, wherein the processor is configured to determine that the object is embedded in the tire by executing further computer instructions to:

- segment the audio data into data blocks whose size is calculated based on an expected length of time of a single rotation of the tire;
- detect an audio anomaly in the audio data based on a degree of correlation between the data blocks; and

determine that the object is embedded in the tire based on the audio anomaly.

19. The computing device of claim **15**, wherein the processor is configured to determine that the object is embedded in the tire by executing further computer instructions to:

- perform a Pearson product-moment coefficient technique on the audio data to detect a time correlated audio anomaly indicative of the object being embedded in the tire.

20. The computing device of claim **15**, wherein the processor is configured to determine that the object is embedded in the tire by executing further computer instructions to:

- employ a trained machine learning model on the audio data to classify the audio data as containing a frequency dependent anomaly indicative of the object being embedded in the tire.

* * * * *