DETECTION OF IMPAIRED DRIVING

TECHNICAL FIELD

[0001] The present disclosure generally relates to detection of impaired driving of vehicles.

BACKGROUND

[0002] A vehicle operator or driver may be impaired for many reasons, including, for example, any one or more of: alcohol intoxication, consumption of drugs such as depressant drugs (e.g., cannabis), stimulants (e.g., methamphetamine, ecstasy, cocaine), sedatives (e.g., benzodiazepine), opioids (e.g., oxycodone, morphine, methadone), MDMA, inhalants, hallucinogens, dissociative anesthetics, narcotic analgesics, heroin, fatigue or sleep deprivation, and exhaustion.

[0003] Current solutions to detecting impaired driving involve in-cabin monitoring solutions within the vehicle. Such solutions may be prone to safety or security problems such as the vehicle operator i.e., driver or other passengers tampering with the monitoring system by covering or obscuring the in-cabin cameras and/or other sensors. Further, such systems suffer from various limitations; for example, they may only be able to monitor a specific target vehicle, and they may be limited to in-vehicle sensing and/or specific actuator elements (e.g., steering wheel vibrations).

[0004] Vision based approaches for detecting impaired driving rely on conventional computer vision techniques for extracting background, detecting and segmenting the desired objects therefrom. However, the input image quality being poor limits the amount of change in vehicle track that can be accurately captured. Such approaches may also be limited in scalability and adaptability to varied environmental conditions, and require extensive pre-processing to achieve desired accuracy levels.

[0005] It is desired to overcome or ameliorate one or more difficulties of the prior art, or to at least provide a useful alternative.

SUMMARY

[0006] According to an aspect of the invention, there is provided a method of detecting impaired driving by a vehicle operator actively operating a vehicle, the method comprising:

detecting an actively-operated vehicle in a monitoring zone; capturing a plurality of images of the detected vehicle using at least one camera; generating a plurality of point cloud frames associated with the detected vehicle using a plurality of ranging sensors; determining a speed and a trajectory of the detected vehicle based on the captured plurality of images and the generated plurality of point cloud frames; and determining an impairment likelihood score based on the determined speed and trajectory for detecting impaired driving.

[0007] According to another aspect of the present invention, there is provided a system for detecting impaired driving by a vehicle operator actively operating a vehicle, the system comprising: a detection module comprising at least one camera and a plurality of ranging sensors, wherein the detection module is configured to: detect an actively-operated vehicle in a monitoring zone; capture a plurality of images of the detected vehicle using the at least one camera; and generate a plurality of point cloud frames associated with the detected vehicle using the plurality of ranging sensors; and an impairment determination module communicatively coupled to the detection module and configured to: determine a speed and a trajectory of the detected vehicle based on the captured plurality of images and the generated plurality of point cloud frames; and determine an impairment likelihood score based on the determined speed and trajectory, for detecting impaired driving.

[0008] The determined trajectory may comprise a lateral position of the detected vehicle.

[0009] Determining the impairment likelihood score may further be based on a type of the detected vehicle. Determining the impairment likelihood score may further be based on one or more parameters associated with external stimuli determined to be present in the monitoring zone.

[0010] The monitoring zone may be determined based on at least one of a type of the at least one camera, a total number of cameras, and a type of the plurality of ranging sensors, a total number of the plurality of ranging sensors.

[0011] The plurality of ranging sensors may comprise a LiDAR sensor and a radar sensor.

[0012] The vehicle operator may be impaired by any one or more of: alcohol consumption, drug use, fatigue.

[0013] The plurality of images captured by the at least one camera may comprise a plurality of in-cabin images of the detected vehicle.

[0014] The plurality of in-cabin images may comprise at least one steep image and at least one shallow image of the detected vehicle.

[0015] A mobile device distraction likelihood score may be determined based on the captured at least one steep image and at least one shallow image of the detected vehicle.

[0016] A seatbelt non-compliance likelihood score may be determined based on the captured at least one steep image and at least one shallow image of the detected vehicle.

[0017] An object of interest likelihood score may be determined based on the captured at least one steep image and at least one shallow image of the detected vehicle.

[0018] The impairment likelihood score may be determined further based on the determined mobile device distraction likelihood score and/or the determined seatbelt non-compliance likelihood score and/or the determined object of interest likelihood score.

[0019] The determined impairment likelihood score may be sent to a communication device for notifying an enforcement officer.

BRIEF DESCRIPTION OF THE DRAWINGS

[0020] Preferred embodiments of the present disclosure are hereafter described, by way of non-limiting example only, with reference to the accompanying drawings, in which:

a. Fig. 1 is a flow diagram of a method of detecting impaired driving by a vehicle operator actively operating a vehicle, in accordance with an embodiment of the present invention;
b. Fig. 2 is a block diagram of a system for implementing the method of Fig. 1, in accordance with an embodiment of the present invention;

c. Fig. 3 is a block diagram of a detection module of the system of Fig. 2, in accordance with an embodiment of the present invention;
d. Fig. 4 is an exemplary image of deployment of the system of Fig. 2 on a transportable trailer unit, in accordance with an embodiment of the present invention;
e. Fig. 5 is an exemplary view showing a monitoring zone and a vehicle path, in accordance with an embodiment of the present invention;
f. Fig. 6 is one of a plurality of images captured by a camera of the system of Fig. 2;
g. Fig. 7 is one of a plurality of point cloud frames associated with the detected vehicle generated using a LiDAR sensor;
h. Fig. 8A illustrates detection of an actively operated vehicle in a monitoring zone, in accordance with an embodiment of the present invention;
i. Fig. 8B illustrates detection of another actively operated vehicle in the monitoring zone, in accordance with an embodiment of the present invention;
j. Fig. 8C is an exemplary processed vehicle list of the detected vehicles of Fig. 8A, 8B generated using a radar sensor, in accordance with an embodiment of the present invention;
k. Fig. 9 is a flow diagram of a method of detecting impaired driving by a vehicle operator actively operating a vehicle including distracted driving and/or driving without proper seatbelt restraint and/or driving while possessing an object of interest associated with alcohol and/or drugs, in accordance with an embodiment of the present invention;
l. Fig. 10 shows examples of images captured i.e. a number plate image (A) and a shallow image (B) taken using a shallow angle camera of an imaging module, a steep image (C) taken using a steep angle camera of the imaging module, and an exemplary 3D vehicle trajectory (D) determined by the detection module of Fig. 3;
m. Fig. 11A, 11B, 11C, 11D show examples of vehicle trajectories when weaving, drifting, swerving, turning with a wide radius respectively;
n. Fig. 12 shows two examples of standard deviation of lateral position of a vehicle;

o. Fig. 13 shows a box plot of time to zero throttle reaction time to variable message sign by a vehicle operator, in accordance with an embodiment of the present invention;
p. Fig. 14A, 14 B show understeering and oversteering of a vehicle;
q. Fig. 15A, 15B and 15C illustrate an exemplary deployment setup of the proposed invention according to an embodiment of the present invention; and
r. Fig. 16 is an exemplary deployment setup of an imaging module, in accordance with an embodiment of the invention.

DETAILED DESCRIPTION

[0021] Embodiments of the present invention include a system and method that are able to automatically detect impaired driving by capturing images of a vehicle in a monitoring zone, using ranging sensors to generate point cloud frames associated with the vehicle, processing the images and the point cloud frames to determine a speed and a trajectory of the vehicle, and using these to determine an impairment likelihood score. Based on the impairment likelihood score, the vehicle can then be intercepted to prevent further driving by an impaired driver, or some other action taken to assess the driver and, if appropriate, issue them with a fine, court summons, require them to undergo drug or alcohol counselling, cancel their driving license, and/or any other action to hopefully prevent possible injury and/or death caused by further impaired driving by the driver.

[0022] Fig. 1 is a flow diagram of a method for detecting impaired driving by a vehicle operator actively operating a vehicle in accordance with an embodiment of the present invention. The method 100 comprises detecting an actively-operated vehicle in a monitoring zone (step 102); capturing a plurality of images of the detected vehicle using at least one camera (step 104); generating a plurality of point cloud frames associated with the detected vehicle using a plurality of ranging sensors (step 106); determining a speed and a trajectory of the detected vehicle based on the captured plurality of images and the generated plurality of point cloud frames (step 108); and determining an impairment likelihood score based on the determined speed and trajectory for detecting impaired driving (step 110).

[0023] The vehicle operator i.e., driver, may be impaired by any one or more of: alcohol intoxication, consumption of drug(s) or fatigue. An actively operated vehicle may be a vehicle that is in motion along a path such as a road.

[0024] Fig. 2 is a block diagram of a system 200 for implementing the method of Fig. 1, in accordance with an embodiment of the present invention. The system may comprise a detection module 202 and an impairment determination module 204. The detection module may comprise at least one camera and a plurality of ranging sensors. The detection module 202 may be configured to detect an actively operated vehicle in a monitoring zone; capture a plurality of images of the detected vehicle using at least one camera; and generate a plurality of point cloud frames associated with the detected vehicle using the plurality of ranging sensors. The impairment determination module 204 may be communicatively coupled to the detection module 202 and may be configured to: determine a speed and a trajectory of the detected vehicle based on the captured plurality of images and the generated plurality of point cloud frames; and determine an impairment likelihood score based on the determined speed and trajectory, for detecting impaired driving.

Detection Module

[0025] The detection module 202 as shown in Fig. 3 may comprise the following sensors: a camera 302, a Light Detection And Ranging (LiDAR) sensor 304 and a RAdio Detecting And Ranging (radar) sensor 306. In a preferred embodiment, the camera 302, the LiDAR sensor 304 and the radar sensor 306 may be attached to a pole e.g. a tripod (preferably substantially vertical) fixed along a side of the path or road, attached to a transportable or moving trailer unit along the roadside, or attached to or mounted on fixed structures such as a gantry, a mast arm or a luminaire. In an embodiment, the height at which the detection module 202 may be deployed can range from 7m to 10m (fixed sites), 7m to 8m (movable sites).

[0026] FIG. 4 illustrates an exemplary deployment setup of the detection module 202 comprising the camera 302 inside a camera enclosure, a LiDAR sensor 304, a radar sensor 306 rigidly mounted on an extendable arm of a transportable trailer unit, and the impairment determination module 204 (not shown). The detection module 202 may additionally comprise one or more infrared (IR) illuminators (Fig. 4 shows two infra-red (IR) illuminators by way of an example). In one embodiment, the IR illuminator may have a power rating between 50W to

200W. By being rigidly attached i.e. in a manner that the camera 302, the LiDAR sensor 304 and the radar sensor 306 do not move with respect to one another in any direction, orientation, or axes, the calibration results, in the form of transformation matrices from all the sensors may be accurate and valid. In an embodiment, the camera 302 in the detection module 202 may be optimized for a scene (an entire field of view) for capturing vehicles in motion for vehicle monitoring and vehicle trajectory determination. Such optimization may be achieved by including a constant on infrared (IR) illuminator to illuminate the scene i.e. monitoring zone and using camera 302 (frame rate may be in the range of 20 to 30fps) to capture all vehicles passing through the monitoring zone such that the vehicle trajectories can be derived with higher precision. In an embodiment, the lens of the camera 302 may have a focal range of 25mm to 60mm. The camera 302 may be configured to continuously capture image frames of vehicles passing through the monitoring zone.

[0027] In some embodiments, the detection module 302 may comprise a plurality of cameras 302, a plurality of LiDAR sensors 304, a plurality of radar sensors 306, the exact numbers being chosen depending on a monitoring zone coverage and/or distance required. In one embodiment, the monitoring distance or length of the monitoring zone is 120m when the detection module 202 comprises a single camera 302, a single LiDAR sensor 304 and a single radar sensor 306. Accordingly, a vehicle travelling at 60km/h would be monitored for 7.2s. In another embodiment, the monitoring distance or length of the monitoring zone may be extended from 120 m to 240 m or more when the detection module 202 comprises two or more LiDAR sensors 304 (e.g., 360° scanning LiDAR sensors) instead of a single LiDAR sensor. In various embodiments, the monitoring zone may be in the range of 120m to 140m in distance and may include a sensor dead zone to accommodate for a zone directly beneath the mounting structure at second deployment site where the sensors in the detection module 202 may not be able to track or detect the vehicle(s).

[0028] It will be appreciated that the camera 302, the LiDAR sensor 304 and the radar 306 may be pre-calibrated such that raw data obtained from the camera, the LiDAR sensor and the radar sensor can be translated between their respective coordinate systems. The pre-calibration/standardization may be performed by taking into consideration one or more of the following factors: manufacturing process, assembly of the sensors. The hardware and component selection of the sensors may be optimized to reduce the impact of time of day and

weather conditions. Standardizing the sensors may reduce or eliminate the impact on data quality due to changes associated with time of day, weather conditions, assembly and/or manufacturing process.

[0029] In one embodiment, the camera 302 may be a high resolution, high frame rate industrial camera (e.g., resolution of 12MP, frame rate of 30 frames per second ("fps")). The camera 302 may comprise one or more filters. The camera 302 may or may not comprise an infrared (IR)-blocking filter. The camera 302 may comprise a narrow band filter applied to a front or rear of the one or more lens. The narrow band filter may let through only the wavelengths of light produced by the respective flash of the one or more flash. This narrow band filter may eliminate the majority of ambient light and/or light produced by the sun. Consequently, resulting image illumination may be controlled more precisely through the camera 302, with minimal impact from sunlight and other environmental sources. The camera 302 may be configured to operate in conjunction with infrared flash units which may provide consistent image quality in all lighting and weather conditions.

[0030] When the one or more flash comprises an 850 nm flash, the one or more filter may let through light only between 700 and 1000nm; 750 and 950nm; 800 and 900nm; 820 and 890nm; 830 and 880nm; 850 and 870nm; or 840 and 860nm. The one or more filter may eliminate about %; about 95% or about 97.5 % of the light normally visible by the camera 302, letting through only the light at the same wavelength as the respective one or more flash.

[0031] The narrow band filter may block all light or substantially all light except that light at or around a particular wavelength. The narrow band filter may comprise a wavelength band of less than 5nm; 5; 10; ;20;25;30;35;40;45;50;55;60;65;70;75;80 85;90;95;100;110;120;130;140; 150; 200; 250; 300; 350; 400; 450; or 500 nm. The narrow band filter may comprise a wavelength band of 5 nm or less; 10 nm or less; 15 nm or less; 20 nm or less; 25nm or less; nm or less; 35nm or less; 40nm or less; 45nm or less; 50 or less; 55nm or less; 60nm or less; 65nm or less; 70nm or less; 75nm or less; 80nm or less; 85nm or less; 90nm or less; nm or less; or 100nm or less.

[0032] In one particular embodiment, the narrow band filter may comprise a Bi850 Near-IR Interference Band pass Filter. In another particular embodiment, narrow band filter may comprise a useful range of 845 to 860nm that is a range of 15 nm.

[0033] The camera 302 may comprise one or more lens such as, a varifocal lens. The camera 302 may comprise a C mount lens or a larger format camera. The camera 302 may comprise a fixed lens such as, an industrial fixed focal length lens. The industrial fixed focal length lens may comprise a rating of 12MP.

[0034] The one or more lens according to any one of the above embodiments may comprise a focal length of 10 to 100mm; 20 to 80 mm or 30 to 60mm. For mobile installation, the focal length may comprise 20 to 50 mm; 25 to 45 mm; or 30 to 40 mm. For mobile installation, the focal length may comprise 35 mm. For fixed installation, the focal length may comprise to 65 mm; 30 to 60 mm or 45 to 55 mm. For fixed installation, the focal length may comprise 50mm. The focal length may be selected to provide tight zoom onto the vehicle for higher resolution and enough width and context to show one or more of: an entire width of the vehicle; most of a lane; and a vehicle number plate.

[0035] According to any one of the above embodiments, the camera 302 may or may not comprise a polarizer.

[0036] In one embodiment of any one of the above embodiments, the camera 302 may comprise one or more flash for illuminating the detected vehicle or a part thereof. The one or more flash may comprise one or more of a 760nm and an 850nm flash. The one or more flash may be capable of firing 10,000 to 100,000; 20,000 to 80,000 or 30,000 to 50,000 times per day at high intensity and short duration. In one embodiment, the one or more flash is capable of firing 40,000 times per day at high intensity and short duration. The one or more flash may comprise one or more light source, the one or more light source may comprise one or more LED light source and/or one or more laser light source. The one or more light source may comprise a narrow-spectrum. The one or more light source may comprise 10 to 1,000; 40 to 500; 300 to 400 light sources. Each of the one or more light sources may comprise an IR LED light source. The one or more light source may be tightly aimed using individual lenses. The individual lenses may be at 1 to 35; 15 to 30; or 20 to degrees. In one embodiment, the individual lenses are at 22 degrees. The one or more flash may comprise one or more capacitor bank to store charge between flashes. The one or more light source when triggered may generate a high intensity of light for a very short duration.

[0037] In one particular embodiment of any one of the above embodiments, the one or more camera comprises a 12MP C Mount camera with a Sony Pregius global shutter sensor.

[0038] According to any one of the above embodiments, the light source may comprise a tightly controlled wavelength. The tight control may comprise a narrow spectral band. The one or more camera filter may exclude all light not within the controlled spectral band. In one embodiment, the tightly controlled wavelength comprises a narrow spectrum light source. The narrow spectrum light source may transmit with a full width at half maximum (FWHM) spectral bandwidth of 5; 10; 20; 25; 30; 35; 40; 45; or 50 nm.

[0039] In a particular embodiment of any one of the above forms, the one or more light source may comprise an Oslon Black, Oslon Black Series 850 nm -80°, SFH 4715AS available from Osram Opto Semiconductors.

[0040] According to any one of the above forms, the spectral bandwidth may be determined at 50% Irel,max full width at half maximum (FWHM).

[0041] The one or more flash may comprise a main flash and a separate offset flash for license plates.

[0042] One or more camera setting may be changed between capture of respective images comprised in the one or more image. The setting may be changed rapidly and/or automatically. The one or more camera setting may comprise exposure time and/or flash intensity.

[0043] According to any one of the above embodiments, an auxiliary camera may be comprised. The auxiliary camera may capture one or more image of a vehicle license plate.

[0044] In one embodiment, the LiDAR sensor 304 may be a high resolution, long-range LiDAR sensor (e.g., 1.5m to 500m detection range, 0.09°x0.08° image-grade resolution). In some embodiments, the LiDAR sensor 304 may have a frame rate between 5 to 20 fps, operate in 1550nm laser wavelength. In an embodiment, the data acquisition rate of the LiDAR sensor 304 may be 10Hz to 20Hz. In an embodiment, the radar 306 may be a high definition radar (e.g., 4D/Ultra High Definition). In some embodiments, the radar 306 may measure range, radial speed, horizontal or azimuth angle, vertical or elevation angle, reflectivity and so on of one or more actively operated vehicles simultaneously. In some embodiments, the radar 306

may be unaffected by weather, temperature and/or lighting conditions. In an embodiment, the radar 306 may operate in 24GHz frequency for detecting actively operated vehicles in single or multiple paths (e.g., 12 lanes). The horizontal or azimuth angle of the radar 306 may range from -25° to +25°, and the vertical or elevation angle may range from -6° to 0°. In an embodiment, the radar 306 may be configured to have horizontal or azimuth angle set to -10° and the vertical or elevation angle set to -2°. In an embodiment, the data acquisition rate of the radar sensor 306 may be 10Hz to 20Hz.

[0045] Some exemplary position- and orientation-related values of the camera 302, LiDAR sensor 304, radar sensor 306 are provided below:

Camera 302	**Value**
Offset from beam centre	2900 mm
Pan Angle (Yaw)	15°
Elevation Angle (Pitch)	-5° tilt of internal camera.

LiDAR sensor 304	**Value**
Offset from beam centre	2050 mm
Pan Angle (Yaw)	15°
Elevation Angle (Pitch)	10°

Radar sensor 306	**Value**
Offset from beam centre	1460 mm

Pan Angle (Yaw)	15°
Elevation Angle (Pitch)	8.5°

[0046] The detection module 202 may perform a 3D multi-object tracking by taking the outputs from the camera 302, LiDAR sensor 304, radar sensor 306 and produce a 3D trajectory for each passing vehicle. The 3D trajectory for each passing vehicle may consist of time series data in the form of (x, y, z, width, length, height, yaw, pitch, roll). An exemplary 3D trajectory dataset may be as follows:

Time (ms)	x (m)	y (m)	z (m)	width (m)	length (m)	height (m)	yaw (rad)	pitch (rad)	roll (rad)
0	-120.00	9.50	-8.00	1.97	4.90	1.80	0.00	0.00	0.00
100	-118.33	9.40	-8.00	1.97	4.90	1.80	0.00	0.00	0.00
200	-116.67	9.40	-8.00	1.97	4.90	1.80	0.00	0.00	0.00
300	-115.00	9.40	-8.00	1.97	4.90	1.80	0.00	0.00	0.00
400	-113.33	9.40	-8.00	1.97	4.90	1.80	0.00	0.00	0.00
500	-111.67	9.40	-8.00	1.97	4.90	1.80	0.00	0.00	0.00
600	-110.00	9.30	-8.00	1.97	4.90	1.80	0.00	0.00	0.00

[0047] The sensors i.e. camera 302, LiDAR sensor 304 and radar sensor 306 may be pre-calibrated by estimating the sensor poses with respect to one or more calibration board(s) and transforming the individual reference frames to a target reference frame.

[0048] The sensors may be configured to continuously capture data such that an actively operated vehicle can be detected upon entering a monitoring zone. Once an actively operated vehicle is detected as having entered a monitoring zone, the camera 302 continuously captures a plurality of images of the detected vehicle, the LiDAR sensor 304 and radar sensor 306 continuously generate point cloud frames associated with the detected vehicle. In some embodiments, the radar sensor 306 may be configured to continuously generate sparse point cloud frames. The sparse point cloud frames may then be post-processed using the metal surface reflections obtained by the radar sensor 306 to generate a processed target list or object list (as shown in Fig. 8C). In other embodiments, the radar sensor 306 may directly generate a processed target or object list.

[0049] Fig. 5 is an exemplary view showing a monitoring zone and a vehicle path, in accordance with an embodiment of the present invention. In some embodiments, the monitoring zone may be sectioned out of the main path or road using obstacles such as traffic cones as shown in Fig. 5. In various embodiments, the monitoring zone may be 120m to 140m in length or distance.

[0050] Fig. 6 is an exemplary image captured by the camera 302 of the system of Fig. 2. It will be appreciated that the combination of camera and camera lens and/or camera configuration used to obtain the image shown in Fig. 6 may be changed to accommodate the monitoring zone or monitoring conditions such as location of monitoring, traffic, time of day or week or year, jurisdictional requirements etc.

[0051] Fig. 7 is an example of one of a plurality of point cloud frames associated with the detected vehicle generated using a LiDAR sensor installed on a transportable trailer unit.

[0052] All collected data (i.e. the captured plurality of images and the generated point cloud frames) are paired with the exact timestamp (time of acquisition) in order to generate a 3D vehicle trajectory in time series.

[0053] Simulator or synthetic dataset may also be used for dataset generation. For example, 3D detection and tracking dataset can be generated using simulators (e.g., carla open source simulator).

[0054] Parts or all of the 3D data collected may then be annotated manually. In some embodiments, the 3D data (i.e. LiDAR sensor point clouds) may be annotated including 3D bounding boxes and corresponding unique vehicle IDs.

[0055] Prior to the detection system being put into operation, the detection system may be configured to detect all passing traffic for a predetermined observation period and establish a baseline driving pattern on the detectable section(s) of the path. Establishing such a baseline may assist with identifying changes in data distribution when the detection system is in operation. The established baseline may then be used as an input to outlier detection in the various vehicle tracking parameters and/or machine-learning algorithms.

Active Lane Detection

[0056] The active lane may be detected by checking if there is active traffic passing by the deployed site. Active traffic may be detected by: radar sensor 306 which may provide a processed object list (as shown in Fig. 8C) and which can be used to determine exact positions of passing traffic or by camera 302 which may provide visual confirmation through an object detection algorithm (e.g., YOLOv5, YOLOv8 algorithms that may be fine-tuned on custom-curated datasets). YOLO (e.g., Ultralytics YOLOv8) may provide a significant advancement in object detection and tracking. This deep learning model may be trained on a vast dataset of images, enabling it to recognize and differentiate objects (vehicles, in this case) in diverse lighting and weather conditions, from various angles, and amidst dense traffic where classical computer vision methods may struggle.

[0057] During the pre-deployment phase, an approved equipment operator can confirm the detected active traffic. If the deployed area or monitoring zone has low traffic volume, the equipment operator may adjust the configuration to add or remove enforced lane(s) or manually adjust and override alignment angle(s) of the sensor(s). If an override alignment angle is provided, the override alignment angle may be used to update configurations of the camera 302, radar sensor 306 and LiDAR sensor 304.

[0058] Prior to detecting active lane(s), a machine learning model may initially be trained using manually marked and labelled images and point clouds that may have been annotated and confirmed by internal or external data annotation team(s). To such a trained ML model, the

following inputs may be provided during operation: (a) image frames obtained from the camera 302, (b) point cloud frames obtained from the LiDAR sensor 304. The detection module 202 may then detect the active lane on the path, boundaries of the active lane on the path on which the target vehicle is moving.

[0059] Fig. 8A illustrates detection of an actively operated vehicle in a monitoring zone. A road may consist of two lanes - lane 1 and lane 2, having a monitoring zone A 702 and monitoring zone B 704 respectively as shown in Fig. 8A. In operation, the detection module 202 may detect an actively operated vehicle 706 travelling on lane 1. The radar 306 may detect the instantaneous speed of the vehicle 706 in the monitoring zone A 702.

[0060] Fig. 8B illustrates detection of another actively operated vehicle. Here too, the road may consist of two lanes - lane 1 and lane 2, having a monitoring zone A 702 and monitoring zone B 704 respectively as shown in Fig. 8B. In operation, the detection module 202 may detect another actively operated vehicle 708 travelling on lane 1. The radar 306 may detect the instantaneous speed of the vehicle 708 in the monitoring zone A 702.

[0061] Fig. 8C is an exemplary processed vehicle list of the detected vehicles of Fig. 8A, 8B generated using the radar 306. Each vehicle will be assigned a unique Object ID/ vehicle ID depending on radar reflection surface (e.g. reflection from the metal surfaces on motorcycles, vehicles). The instantaneous speed of the detected vehicles will be determined by the radar 306 while the active lane and monitoring zone will be determined by a combination of the camera 302, LiDAR sensor 304 and radar sensor 306. A class of the vehicle may be selected from a predetermined list of classes. In an exemplary embodiment, the vehicle class may be: 0 - undefined, 1 – pedestrian, 2 – Bicycle, 3 – Motorbike, 4 – car, 5 – transporter, 6 – short truck, 7 – long truck.

[0062] It will be appreciated that the number of lanes that may be tracked and/or monitored for detecting vehicles may be based on where the system 200 is deployed. In an embodiment, when system 200 is mounted on a transportable trailer unit, the detection module 202 may track and/or monitor two lanes simultaneously. In another embodiment, if the system 200 is deployed at a fixed site, the detection module 202 may track and/or monitor more lanes simultaneously (e.g. 4 or 5 lanes). It will further be appreciated that the detection module 202 may track

movement of vehicles "across" lanes. For instance, if a target vehicle is weaving between Lane 1 to 2 to 3 in an S-like pattern, such movement may be captured for analysis.

Trajectory of detected actively operated vehicle

[0063] The detection module 202 uses the captured plurality of images and generated point cloud frames to determine the active lane, active lane boundaries. Once the active lane and active lane boundaries are determined, a 3D vehicle trajectory in the form of time series positional data may be determined. Vehicle trajectory is a time series of vehicle positions.

[0064] Lateral position measurement can capture vehicle operator behaviour such as weaving, drifting, swerving and turning. Fig. 11A, 11B, 11C, 11D show examples of vehicle trajectories when weaving, drifting, swerving, turning with a wide radius respectively.

[0065] A mean lateral position (MLP) of the detected vehicle may be calculated for the entire monitoring zone length e.g., 120m.

[0066] If 'X' represents the lateral position of the detected vehicle with a mean value μ, calculated across all valid samples, the mean lateral position (MLP) may be calculated as follows:

$$MLP\,[X] = \mu \tag{1}$$

[0067] A standard deviation of mean lateral position (SDLP) may then be calculated across all samples obtained over the monitoring zone length as follows:

$$SDLP = \sqrt{MLP\,[(X-\mu)^2]} \tag{2}$$

[0068] SDLP may serve as a robust and sensitive indicator of vehicle operator impairment. Fig. 12 shows two examples of standard deviation in lateral position of a vehicle

[0069] The SDLP value may then be normalized to a [0.0, 1.0] scale using min-max normalization. This adjustment may account for the expected range of SDLP values under normal driving conditions. The normalized SDLP score, S_{SDLP}, may be calculated as follows:

$$S_{\text{SDLP}} = \frac{\text{SDLP} - \min(\text{SDLP})}{\max(\text{SDLP}) - \min(\text{SDLP})} \tag{3}$$

[0070] Here, min(SDLP) and max(SDLP) represent the minimum and maximum observed SDLP values in a calibration dataset, encompassing a broad spectrum of driving behaviours.

Speed of detected actively operated vehicle

[0071] Vehicle speed may be monitored using the camera 302, the LiDAR sensor 304, the radar 306 or a combination of any of these sensors. The camera 302 and the LiDAR sensor 304 may determine the speed of the detected vehicle based on a change in distance over time. The radar 306 may determine the speed of the detected vehicle based on Doppler Effect.

[0072] The system 200 may use a sensor fusion-based approach, utilizing a 3D tracking algorithm to merge data output from the camera 302, LiDAR sensor 304, and Radar sensor 306 sensors. This approach may improve upon the standard radar-based speed measurement by also taking into consideration the inputs from the camera 302 and LiDAR sensor 304 sensors thereby eliminating false speed readings caused by metal reflections and other inaccuracies in challenging environments.

[0073] A mean speed of the detected vehicle may be calculated for the entire monitoring zone length e.g., 120m.

[0074] If 'X' represents the speed of the detected vehicle with a mean value μ, calculated across all valid samples, the mean speed (MS) may be calculated as follows:

$$\text{MS}\,[X] = \mu \tag{4}$$

[0075] A standard deviation of speed (SDS) may then be calculated across all samples obtained over the monitoring zone length as follows:

$$\mathrm{SDS} = \sqrt{\mathrm{MS}[(X-\mu)^2]} \tag{5}$$

[0076] The SDS value may then be normalized to a [0.0, 1.0] scale using min-max normalization. This adjustment may account for the expected range of SDS values under normal driving conditions. The normalized SDS score, S_{SDS}, may be calculated as follows:

$$S_{SDS} = \frac{\mathrm{SDS} - \min(\mathrm{SDS})}{\max(\mathrm{SDS}) - \min(\mathrm{SDS})} \tag{6}$$

[0077] Here too, min(SDLP) and max(SDLP) represent the minimum and maximum observed SDLP values in a calibration dataset, encompassing a broad spectrum of driving behaviours.

[0078] It will be appreciated that while use of vehicle speed in determining the impairment likelihood score is being discussed in this disclosure, other vehicle speed-related parameters such as, but not limited to, vehicle velocity or acceleration or the full 2D/3D vehicle trajectory may alternatively be used in determining the impairment likelihood score. In an embodiment, a birds-eye-view machine learning model may be used to predict the vehicle trajectory on a 2D plane in place of 3D trajectory that includes the vehicle height, vehicle roll, vehicle pitch.

Impairment Determination

[0079] A weight factor may be assigned to each of the normalized standard deviation in lateral position S_{SDLP} and normalized standard deviation in speed S_{SDS} depending on its individual importance in indicating the impairment of the vehicle operator.

[0080] In an exemplary embodiment, a weight factor W_{SDLP} is given to the vehicle trajectory and a weight factor W_{SDS} is given to the vehicle speed and the impairment likelihood score (O) may be calculated as a weighted average as follows:

$$W_{SDLP}.S_{SDLP} + W_{SDS}.S_{SDS}$$

$$O = \frac{}{W_{SDLP} + W_{SDS}} \quad (7)$$

[0081] The overall impairment likelihood score may be determined using machine learning using deep neural networks as follows. All raw data and corresponding collected parameters discussed above may be mapped to collected ground truth labels. Inputs to the deep neural network may comprise raw sensor data (including camera image frames, LiDAR sensor point cloud frames, radar sensor sparse point cloud frames), processed data and parameters (including vehicle type, vehicle velocity, vehicle trajectories), in-cabin images (shallow image and steep image), baseline data of passing vehicles (from observation phase). Output of the deep learning network may comprise a classification likelihood (ground truth labels used for training the neural network may be based on alcohol and/or drug test results provided by partnered government agencies). The classification likelihood would indicate whether or not the vehicle operator is impaired. The dataset may consist of N samples of (input, output) pairs which may be split into training, validation and testing sets. Using the training set, the deep neural network(s) are trained to predict the impairment likelihood based on all provided inputs. The validation set may be used to fine-tune the neural network parameters. The testing set may be used to measure how well the neural network generalizes to data outside of the training set.

[0082] The overall impairment likelihood score may also be determined using non-machine-learning based approach i.e., an outlier detection method wherein using the same set of input data mentioned above, the outlier detection method may be tuned based on weighting of key parameters such as: vehicle type, vehicle trajectory relative to marked road boundaries, vehicle velocity (including other related parameters such as standard deviation of lateral position, acceleration), vehicle operator attentiveness (using upstream in-cabin imagery which checks if the vehicle operator is distracted by using mobile device(s). The tuning or weighting may be affected by the measured baseline of passing traffic using an initial observation phase. Baseline may be based on observing all passing traffic and establishing an average driving pattern.

Vehicle Type

[0083] The actively operated vehicle may be a heavy vehicle (e.g., a truck) or a normal vehicle (e.g., a car), a two-wheeler (e.g., a motorbike) etc. Change in speed and trajectory of a vehicle

may be impacted differently based on the type of vehicle. In one example, a commercial heavy vehicle may be less responsive to input by vehicle operator in comparison to a car and therefore the speed and trajectory may be less indicative of vehicle operator impairment in some scenarios. Accordingly, the weight factor W_{SDLP} and W_{SDS} may be chosen depending on the type of vehicle in determining the impairment likelihood score.

Vehicle Operator Parameters

[0084] The operator i.e. driver of the actively operated vehicle may be monitored to determine reaction time. Reaction time may be any one or more of the following: reaction time of the vehicle operator to posted speed sign or variable message sign, stop sign, any advisory signs, changes to speed limits, traffic or other lights, speed bumps, sudden change in path layout e.g., sharp turns, obstacles, traffic cones, rocks or boulders, faller trees, external stimuli (e.g., an object or person(s) crossing in front of the vehicle). Reaction time may be responsive to deceleration (lifting foot off the accelerator paddle and pressing on brake in response to an external stimulus) or acceleration (lifting foot off the brake paddle and pressing on accelerator paddle in response to external stimulus) of the vehicle by the vehicle operator. Fig. 13 shows a box plot of time to zero throttle (in seconds) reaction time to variable message sign (VMS) by a vehicle operator.

[0085] In an exemplary embodiment, if a vehicle is moving through an urban area at 30 km/h and a pedestrian suddenly runs across the road unexpectedly, the situation would require immediate response from the vehicle operator to prevent a potential collision. In such a scenario, the camera 302 and LiDAR 304 sensors may detect the unexpected pedestrian (external stimuli). The sensors may determine information related to the pedestrian's speed, trajectory, and distance from the vehicle, track the pedestrian's movement, and predict their path and speed. The sensors may assess the risk of potential collision by calculating the time to collision (TTC) if the vehicle continues on its current path without adjusting speed or direction i.e. if the vehicle speed and/or trajectory is not changed by the vehicle operator.

[0086] Accordingly, the response of the vehicle and/or the vehicle operator is simultaneously analysed. If the vehicle and/or the vehicle operator initiates a sudden deceleration or steering action, such changes may be assessed in the context of the external stimulus. A swift and appropriate response (e.g., braking or evasive manoeuvring within safety limits which would

imply change in the speed and/or trajectory) would indicate proper situational awareness and low impairment. If the vehicle's response is delayed, inadequate, or overly aggressive, the impairment likelihood score would increase, signalling potential impairment or system failure.

[0087] The response of the vehicle and/or the vehicle operator, combined with the real-time analysis and intervention capabilities of the system 200, may play a crucial role in safely managing such unexpected situations. Such incidents may contribute to the continuous learning of the system 200, enhancing its predictive models and response strategies for future encounters.

[0088] An impaired driver may tend to undershoot or overshoot lane changes and/or turns. Undershoot (or understeer) is when the vehicle operator turns less sharply than is intended/required to safely navigate a turn (as shown in Fig. 14A). Overshoot (or oversteer) is when the vehicle operator turns more sharply than is intended/required to safely navigate a turn (as shown in Fig. 14B).

Communication device

[0089] The notification may be sent in the form of a message including the impairment likelihood score to one or more computing devices or communication devices of law enforcement unit(s). An enforcement officer e.g., a traffic police officer may use the notification to decide whether or not to pull over the detected/flagged actively operated vehicle for further testing e.g., when the impairment likelihood score received at the computing device or the communication device exceeds a predetermined threshold. The predetermined threshold may be set based on the jurisdiction and/or operational capability of the system 200. A higher threshold may provide lesser false positives at the cost of obtaining fewer data samples. Further testing may include one or more of: a random breath test (RBT) operation using a breathalyzer for checking if the vehicle operator has consumed alcohol, a traffic stop and test using a random drug test (RDT) machine for checking if the vehicle operator has taken drugs. In case of fatigued drivers, the enforcement officer may decide to pull over the detected actively operated vehicle to request the vehicle operator to take rest before proceeding with their journey.

[0090] Fig. 9 is a flow diagram of an alternative embodiment of the method of detecting impaired driving by a vehicle operator actively operating a vehicle. The method 900 includes

steps additional to the steps of method 100, for determining distracted driving and/or driving without proper seatbelt restraint and/or presence of object(s) of interest as will be explained below.

[0091] In method 900, a plurality of in-cabin images of the detected vehicle are captured (step 902) based on which one or more of a mobile device (e.g., mobile phone, tablet etc.) distraction likelihood score is determined (step 904), a seatbelt non-compliance likelihood score is determined (step 906) and an object of interest likelihood score is determined (step 908). One or more of the determined mobile device distraction likelihood score, the determined seatbelt non-compliance likelihood score and the determined object of interest likelihood score are then factored into determining the impairment likelihood score (step 910).

[0092] In this embodiment, the mobile device and/or seatbelt non-compliance and/or object of interest detection cameras may be deployed upstream and/or co-located with the proposed impaired driving detection system 200.

[0093] In an example embodiment, if the vehicle operator of the detected vehicle is confirmed to be distracted by a mobile device, then the impairment likelihood score would be impacted. This can also help reduce the false positive rate by significantly discounting the cases where target vehicles were confirmed distracted cases.

[0094] The mobile device and/or seatbelt non-compliance and/or object of interest detection cameras may further comprise a shallow angle camera and a steep angle camera. The shallow angle camera may have settings configured to capture a first image containing the number plate of the detected vehicle and a second image corresponding to a shallow in-cabin image of the detected vehicle. For capturing the second image, the camera settings of the shallow angle camera may be optimized such that the vehicle operator and front seat passenger are captured in the second image. In an example, the shallow angle camera may be tilted downwards to around 10° below the horizontal axis. In various embodiments, the shallow angle camera may be tilted downwards at an angle between 5° to 15° below the horizontal axis. The steep angle camera may have settings configured to capture a third image corresponding to a steep in-cabin image of the detected vehicle. For capturing the third image, the camera settings of the steep angle camera may be optimized such that the vehicle operator and front seat passenger are captured in the third image. In an example, the steep angle camera may be tilted downwards

to around 30° to 35° below the horizontal axis. In various embodiments, the steep angle camera may be tilted downwards at an angle between 25° to 35° below the horizontal axis. It will be appreciated that the exact tilt angles for the shallow angle and steep angle cameras may be chosen depending on the deployment site. Operation of the shallow and steep angle cameras may be optimized by using a flash illuminator with band pass filter as discussed above whereby the in-cabin images can be captured using short pulses of infrared light. Examples of an image containing number plate of a vehicle, a shallow angle image and a steep angle image, as well as an exemplary 3D vehicle trajectory determined by the camera 302, Lidar sensor 304 and radar sensor 306 of the detection module 202 are shown in Fig. 10 ((A), (B), (C), (D) respectively). In an embodiment, the mobile device and/or seatbelt non-compliance and/or object of interest detection cameras may have a frame rate of 8fps and a focal range of 35mm to 50mm. In various embodiments, the monitoring angle and zoom level of the shallow angle camera and the steep angle camera may be based on a mounting height and/or type of lens used.

[0095] In an embodiment, the impairment likelihood score may be calculated as follows: a mobile device distraction likelihood score S_{device} between [0.0, 1.0] may be generated by an image classifier based on the images captured by the shallow angle camera and the steep angle camera, where a score of 0.0 may imply that the vehicle operator is attentive and not distracted, while a score of 1.0 may imply that the vehicle operator is distracted. A weight factor W_{device} is assigned to the mobile device distraction likelihood score based on its relevance in the deployment site, time, or other factors and the overall impairment likelihood score (O) may be calculated as a weighted average as follows:

$$O = \frac{W_{SDLP}.S_{SDLP} + W_{SDS}.S_{SDS} + W_{device}.S_{device}}{W_{SDLP} + W_{SDS} + W_{device}} \quad (8)$$

[0096] In a further embodiment, a seatbelt non-compliance likelihood score $S_{seatbelt}$ may be generated by an image classifier based on the images captured by the shallow angle camera and the steep angle camera, where a score of 0.0 may imply that the vehicle operator complies with the seatbelt requirements i.e. has worn the seatbelt correctly, while a score of 1.0 may imply that the vehicle operator has not complied with the seatbelt requirements i.e. has not

worn the seatbelt correctly. A weight factor $W_{seatbelt}$ is assigned to the seatbelt non-compliance likelihood score based on its relevance in the deployment site, time, or other factors and the overall impairment likelihood score (O) may be calculated as a weighted average as follows:

$$O = \frac{W_{SDLP}.S_{SDLP} + W_{SDS}.S_{SDS} + W_{device}.S_{device} + W_{seatbelt}.S_{seatbelt}}{W_{SDLP} + W_{SDS} + W_{device} + W_{seatbelt}} \quad (9)$$

[0097] In yet another embodiment, an object of interest (e.g., alcohol bottle, laughing gas canister, pipes, bongs rolling papers, syringes or other objects used for the distribution and/or consumption of alcohol or drugs) likelihood score S_{obj} may be generated by an object detector and/or image classifier model based on the images captured by the shallow angle camera and the steep angle camera, where a score of 0.0 may imply that there is no object of interest present in the detected vehicle, while a score of 1.0 may imply that an object of interest is present in the vehicle. A weight factor W_{obj} is assigned to the object of interest likelihood score based on its relevance in the deployment site, time, or other factors and the overall impairment likelihood score (O) may be calculated as a weighted average as follows:

$$O = \frac{W_{SDLP}.S_{SDLP} + W_{SDS}.S_{SDS} + W_{device}.S_{device} + W_{seatbelt}.S_{seatbelt} + W_{obj}.S_{obj}}{W_{SDLP} + W_{SDS} + W_{device} + W_{seatbelt} + W_{obj}} \quad (10)$$

[0098] The weight factors (W_{SDLP}, W_{SDS}, W_{device}, $W_{seatbelt}$, W_{obj}) may be adjusted using calibration datasets to fine-tune the machine learning model to ensure that the overall impairment likelihood score (O) reflects the likelihood of impairment of the vehicle operator more accurately.

[0099] The updated model may be validated with a separate test dataset, assessing the improvement in predictive performance metrics such as accuracy, precision, recall, and the Receiver Operating Characteristic Curve (ROC) - Area under ROC Curve (AUC) score.

[00100] Note the aforementioned "Validation" step may be possible only if there is a fully labelled dataset with a sizable amount of true-positive samples (e.g. 100 true positive samples). Further, the number of true negative samples may be many times more (if not many magnitudes more) than the true positive samples collected.

Addition to Fully Labelled Dataset and Continuous Model Improvement

[00101] Assuming the proposed method is continuously trialed and deployed, and there are partnered agencies who intercept and apply drug and/or alcohol tests to some/all vehicles downstream, then there may be new data samples available which can be annotated and added to the Fully Labelled Dataset.

[00102] Such a Fully Labelled Dataset may be an extension of the Multi-Object 3D Tracking Dataset discussed above. A Fully Labelled Dataset may consist of the following: a first image used to determine number plate of the detected vehicle, a second image (shallow in-cabin image), a third image (steep in-cabin image), raw images from scene camera 302 which captures vehicle movement, targets detected by the radar sensor 306, LiDAR sensor 304 point clouds, Ground Truth Labels, binary output on whether the vehicle was tested for alcohol (Yes/no, if tested, the BAC measurement would be recorded (Decimal Value), and if tested, binary output on whether the BAC measurement exceeded legal limit in the deployed jurisdiction (Yes/no), binary output on whether the vehicle was tested for drug (Yes/no, if tested, the specific drug test result (No Drug Detected / Cannabis / Methamphetamine / Cocaine / Ecstasy ... and so on).

[00103] The full list of drugs tested can vary between different states/jurisdictions. The output data listed above may be supplied by partnered agencies, bundled together with the number plate of the detected and tested vehicle. The input and output data are matched up using the recorded vehicle number plate obtained from the image captured by the camera 302 and the recorded vehicle number plate from the partnered agency.

Transition to Larger Models (If Sufficiently Large Dataset Available)

[00104] The aforementioned approach may be used to continuously collect additional data and fine-tune existing model parameters. However, when the dataset has true positive

samples in the order of thousands to tens of thousands, it may be possible to further improve model performance by training a neural network to directly predict the overall impairment likelihood score using data available. The inputs may comprise vehicle 3D trajectory, in the form of time series positional data, mobile device distraction likelihood score, object of interest likelihood score. The output may comprise the overall impairment likelihood score.

[00105] Fig. 15A, 15B and 15C illustrate an exemplary deployment setup of the proposed invention according to an embodiment of the present invention. An imaging module 1502 comprising at least one shallow angle camera, a steep angle camera and a radar may be deployed at a first deployment site to monitor all approaching vehicles for: mobile device distraction, seatbelt compliance, objects of interest as discussed above. In one embodiment, the imaging module 1502 may monitor vehicles within a 50m approach distance to be tracked by the radar as well as to capture appropriate shallow angle and steep angle images as discussed above. The imaging module 1502 may be mounted on a mounting structure which may be a fixed structure e.g., a pole, tripod etc. or on a transportable trailer unit as shown in Fig. 16. The radar sensor (not shown) in the imaging module 1502 may be configured to determine point speed and/or average speed of the approaching vehicles to determine point speed and/or average speed non-compliance. A score for the point speed non-compliance and/or average speed non-compliance may be standard deviation of speed which may then be used to calculate the overall impairment likelihood score as discussed above. The detection module 202 may be deployed at a second deployment site downstream of the first deployment site to detect and monitor all passing vehicles and extract 3D vehicle trajectories as discussed above. In an embodiment, the distance between the first and second deployment sites may be 140m.

[00106] Fig. 16 is an exemplary deployment setup of the imaging module 1502 on a transportable trailer unit, in accordance with an embodiment of the invention and comprises a pair of steep angle cameras (one for each monitored lane) mounted at a height on the trailer, a shallow angle camera mounted at a height lower than the steep angle cameras, and a radar sensor (not shown).

[00107] For transportable unit operation, when a vehicle approaches the imaging module 1502 on one of the two monitored lanes available, the imaging module 1502 may track the approaching target vehicle using the radar sensor. The shallow angle camera is then triggered

to capture the number plate of the vehicle and a shallow angle image of the vehicle including the vehicle operator and the front seat passenger. The steep angle is triggered to capture the steep image of the vehicle including the vehicle operator and the front row passenger. Once the images are captured, the number plate is processed using a computing device onboard or near the mounting structure and the mobile device distraction likelihood score, seatbelt compliance likelihood score, object of interest likelihood score are determined as discussed above. The captured images and computed scores may be uploaded to a cloud hosted storage and database (e.g., AWS, Microsft Azure, Google Cloud), and can be available for use by the detection module 202 at the second deployment site downstream.

[00108] The detection of a phone or mobile device, seatbelt compliance and/or object of interest ("offences") may comprise searching the one or more image for the phone or mobile device, seatbelt and/or object of interest. The search may comprise a neural network or artificial neural network such as a deep neural network or a deep convolutional neural network. The search may be of an entire image. If a phone or mobile device, improper seatbelt use and/or object of interest is detected, based on its confidence score, then one or more images may be sent for further review. The further review may be by a person. One or more images may be excluded if a detected phone or mobile device is associated with a passenger and not a vehicle operator or driver. One or more image may be automatically excluded if the phone or mobile device is detected in a holder and not grasped by hand.

[00109] The captured and/or received image may comprise a view of the complete front of the vehicle and the lane the vehicle is in.

[00110] The neural network may comprise an object detection system. The object detection system may use a neural network based YOLO (you only look once) real time object detection architecture.

[00111] The neural network may comprise an image classifier. The image classifier may comprise a neural network based VGG (Visual Geometry Group) classifier. The classifier may receive the one or more cropped image. The classifier may be pre-trained. The pre-training may comprise training on a data set such as, the 1,000 class ImageNet set. The output model of the network may be modified to identify only two classes: received cropped images containing illegal phone or mobile device use and/or improper seatbelt use and/or object of

interest, i.e. positive, and received cropped images not containing illegal phone or mobile device use and/or improper seatbelt use and/or object of interest i.e. negative. The classifier may have been trained by fine tuning on a training set of example images. The example images may comprise, more than one hundred, more than one thousand, more than a hundred thousand, or more than one million example images such as, cropped images. The training set may comprise positive and negative labelled images. When the received cropped image of the vehicle operator (e.g, for mobile device use) is fed into the classifier network, the analysis may determine as output a confidence that the image is positive for phone or mobile device use and/or improper seatbelt use and/or object of interest. The determined confidence may comprise for example a confidence output of 0% is very unlikely to show phone or mobile device use and/or improper seatbelt use and/or object of interest and a confidence output of 100% is very likely to show phone or mobile device use and/or improper seatbelt use and/or object of interest. A threshold may be used to limit which images are deemed positive and which are deemed negative. The threshold may comprise a threshold value which may be dynamically adjusted so that a top margin of images are deemed positive and sent on for manual review. The top margin may comprise a top 5; 10; 20; or 25% of received cropped images with respected to determined confidence.

[00112] According to any one of the above forms, the analysis may comprise a real time object detector. The real time object detector may identify the driver by the presence of a steering wheel and the human behind that wheel. The training set may comprise images from various angles optionally, both vertical and horizontal angle variations. The training set may comprise images of many different vehicle types, and optionally images of both left hand drive vehicles and right hand drive vehicles. The analysis may comprise a positive identification of the driver and ignore passengers. The real time object detector may output a location of the driver in the image and optionally a confidence score that the driver has been found.

[00113] According to any one of the above embodiments, when the vehicle operator has been located in the received image, the received image may be cropped.

[00114] The standardised one or more cropped image may be supplied to the image classifier. The image classifier may receive the standardised one or more cropped image. The image classifier may have been trained on a training set comprising a plurality of vehicle

operator cropped image examples. The training set contains labelled images of vehicle operators/drivers illegally touching phones or mobile device, non-illegally touching phones or mobile device, and not touching phones or mobile devices at all. When the image of the driver is received by the image classifier, an output comprising a most likely category and a confidence score of making that determination may be provided.

[00115] The classification may comprise determining the location of a driver in the one or more image. The determination of the location may comprise steering wheel detection and/or person detection. After determination the one or more image may be cropped to generate an image showing only the driver and the driver's immediate surrounds. The immediate surround may comprise a driver's wingspan. The cropped image may then be provided for further review. The further review may be by a person.

[00116] According to any one of the above embodiments, the detection may comprise a confidence factor or threshold. In one embodiment, only the one or more captured images above a threshold are provided. The confidence threshold may be adjusted. The threshold may comprise a setting of 0% of offences missed, 100% of images processed manually to 100% of offences missed, 0% of images processed manually. In one embodiment, the threshold comprises 5% missed for 10% manual processing. Once the target vehicle enters the monitoring zone (e.g., shown as the 120m to 140m zone in Fig. 15B), the detection module 202 detects the vehicle as discussed above. The monitoring zone may be set to begin at a distance of 20 to 100m (buffer zone) from the first deployment site. The sensors (camera 302, LiDAR sensor 304, radar sensor 306) mounted at a mounting structure (e.g., fixed structure like a pole, tripod etc., or movable e.g., transportable trailer unit as shown in Fig. 11) may continuously capture data including the number plate, steep and shallow images, LiDAR sensor and radar sensor point clouds. Using the multi-object 3D tracking algorithm discussed above, the detection module 202 may continuously and incrementally compute the 3D trajectories of all detected objects (including but not limited to vehicles, cyclists and pedestrians). Pedestrian(s) and cyclist(s) trajectories may be used to extrapolate unexpected external stimuli such as a pedestrian crossing the road. Vehicle 3D trajectory may be completed when the target vehicle moves outside of the monitored zone. A computing device onboard or near the mounting structure on which the system 200 is mounted may then compute the speed and trajectory of the target vehicle as discussed above. The impairment determination module may

determine the overall impairment likelihood score as discussed above. It will be appreciated that the computing device and the impairment determination module may be a single computing module or separate computing modules and may either be physically at the second deployment site or remote from the first and/or second deployment site. System 200 may access the results determined by the image module and stored in the cloud-hosted storage to match the determined scores using the number plate. If the calculated overall impairment likelihood score is determined to be greater than or equal to a predetermined threshold, a notification (which may be in the form of an evidence package) may be sent out to downstream staff members, police officers or the like (who may be stationed 200m or more (variable based on available spacing)) downstream from the second deployment site as shown in Fig. 15C) who can then stop the target vehicle for further drug and/or alcohol testing. If the calculated overall impairment likelihood score is determined to be less than a predetermined threshold, no action may be taken. The evidence package may include one or more of the number plate of the detected vehicle, overall impairment likelihood score, vehicle trajectory (which may be overlaid on top of the baseline trajectory), steep and shallow images captured by one or both of the imaging module 1502 and the detection module 202 (which may be overlaid with annotations for mobile device causing distraction to the vehicle operator, improper seatbelt use, or objects of interest detected).

[00117] In an embodiment, if the detected vehicle is intercepted by an enforcement officer, a feedback may be provided from the communication device and/or one or more tripod mounted cameras (as shown in Fig. 15C) associated with the enforcement officer to the system 200. In an exemplary embodiment, the tripod mounted camera(a) may be deployed at 200m or more downstream of the second deployment site. Each interception may include alcohol test result (blood alcohol content) and lab verified drug test result (positive/negative, and the indicated drug). The blood alcohol content (BAC) level, measured in 0.10 g of alcohol for every 100 mL of blood, in Australian jurisdiction for example, if BAC > 0.05%, then positive, else negative). Such feedback may be used by the system 200 to be continuously improved/re-trained.

Curation of dataset for supervised learning of vehicle trajectories

[00118] A high quality, fully labeled dataset may need to be curated for training the neural network for 2D and 3D object detection. The dataset may comprise camera image frames, LiDAR sensor point cloud frames, Radar raw and/or processed targets, ground truth labels provided through a partnered government agency, which can include: random breath test result, in the form of BAC measurement (e.g., in grams of alcohol per 100 ml blood), random drug test result (e.g., positive or negative, and drug identified). After initial data collection, each passing vehicle may be assigned a unique vehicle ID which may then correspond to one of the following categories and sub-categories: True positive (i.e., passing vehicle is tested and confirmed to be impaired by drug and/or alcohol), True negative (i.e., passing vehicle is tested and confirmed to be not impaired), Not assessed (i.e., passing vehicle is not assessed due to administrative reason (e.g. high traffic flow, insufficient testing stations etc.).

Possible Advantages

[00119] One or more embodiments of the present invention may monitor a plurality of actively operated vehicles in the monitoring zone at the same time, provide better safety or security since the sensors are placed externally (mounted on transportable trailer unit, gantry etc.) to the detected vehicles and cannot therefore be tampered with. In some embodiments, the system 200 may be paired with external road elements such as Variable Message Signs (VMS) , traffic cones to measure reaction time(s) of the vehicle operator that can be incorporated into determining the impairment likelihood score. More than one system 200 may be linked to one another to create a networked impairment detection system for monitoring long stretches of roads or covering a large area. The impairment detection time can be reduced. The impairment detection may be optimized.

[00120] In this specification, the terms "comprises", "comprising" or similar terms are intended to mean a non-exclusive inclusion, such that an apparatus that comprises a list of elements does not include those elements solely, but may well include other elements not listed.

[00121] Throughout the specification, the aim has been to describe the invention without limiting the invention to any one embodiment or specific collection of features. Persons skilled in the relevant art may realize variations from the specific embodiments that will nonetheless fall within the scope of the invention.

CLAIMS:

1. A method of detecting impaired driving by a vehicle operator actively operating a vehicle, the method comprising:

detecting an actively-operated vehicle in a monitoring zone;

capturing a plurality of images of the detected vehicle using at least one camera;

generating a plurality of point cloud frames associated with the detected vehicle using a plurality of ranging sensors;

determining a speed and a trajectory of the detected vehicle based on the captured plurality of images and the generated plurality of point cloud frames; and

determining an impairment likelihood score based on the determined speed and trajectory for detecting impaired driving.

2. The method of claim 1, wherein the determined trajectory comprises a lateral position of the detected vehicle.

3. The method of claim 1, wherein determining the impairment likelihood score is further based on a type of the detected vehicle.

4. The method of claim 1, wherein determining the impairment likelihood score is further based on one or more parameters associated with an external stimuli determined to be present in the monitoring zone.

5. The method of claim 1, wherein the monitoring zone is determined based on at least one of a type of the at least one camera, a total number of cameras, and a type of the plurality of ranging sensors, a total number of the plurality of ranging sensors.

6. The method of claim 1, wherein the plurality of ranging sensors comprises a LiDAR sensor and a radar sensor.

7. The method of claim 1, wherein the vehicle operator is impaired by any one or more of: alcohol consumption, drug use, fatigue.

8. The method of claim 1, wherein the plurality of images captured by the at least one camera comprises a plurality of in-cabin images of the detected vehicle.

9. The method of claim 8, wherein the plurality of in-cabin images comprises at least one steep image and at least one shallow image of the detected vehicle.

10. The method of claim 9, further comprising determining a mobile device distraction likelihood score based on the captured at least one steep image and at least one shallow image of the detected vehicle.

11. The method of claim 9, further comprising determining a seatbelt non-compliance likelihood score based on the captured at least one steep image and at least one shallow image of the detected vehicle.

12. The method of claim 9, further comprising determining an object of interest likelihood score based on the captured at least one steep image and at least one shallow image of the detected vehicle.

13. The method of claim 9, 10 or 11, wherein determination of the impairment likelihood score is further based on the determined mobile device distraction likelihood score and/or the determined seatbelt non-compliance likelihood score and/or the determined object of interest likelihood score.

14. The method of claim 1 further comprising sending the determined impairment likelihood score to a communication device for notifying an enforcement officer.

15. A system for detecting impaired driving by a vehicle operator actively operating a vehicle, the system comprising:

a detection module comprising at least one camera and a plurality of ranging sensors, wherein the detection module is configured to:

detect an actively-operated vehicle in a monitoring zone;

capture a plurality of images of the detected vehicle using the at least one camera; and

generate a plurality of point cloud frames associated with the detected vehicle using the plurality of ranging sensors; and

an impairment determination module communicatively coupled to the detection module and configured to:

determine a speed and a trajectory of the detected vehicle based on the captured plurality of images and the generated plurality of point cloud frames; and

determine an impairment likelihood score based on the determined speed and trajectory, for detecting impaired driving.

16. The system of claim 15, wherein the determined trajectory comprises a lateral position of the detected vehicle.

17. The system of claim 15, wherein the impairment determination module is configured to determine the impairment likelihood score further based on a type of the detected vehicle.

18. The system of claim 15, wherein the impairment determination module is configured to determine the impairment likelihood score further based on one or more parameters associated with an external stimuli determined to be present in the monitoring zone.

19. The system of claim 15, wherein the monitoring zone is determined based on at least one of a type of the at least one camera, a total number of cameras, and a type of the plurality of ranging sensors, a total number of the plurality of ranging sensors.

20. The system of claim 15, wherein the plurality of ranging sensors comprises a LiDAR sensor and a radar sensor.

21. The system of claim 15, wherein the vehicle operator is impaired by any one or more of: alcohol consumption, drug use, fatigue.

22. The system of claim 15, wherein the plurality of images captured by the at least one camera comprises a plurality of in-cabin images of the detected vehicle.

23. The system of claim 22, wherein the plurality of in-cabin images comprises at least one steep image and at least one shallow image of the detected vehicle.

24. The system of claim 23, wherein the impairment determination module is further configured to determine a mobile device distraction likelihood score based on the captured at least one steep image and at least one shallow image of the detected vehicle.

25. The system of claim 23, wherein the impairment determination module is further configured to determine a seatbelt non-compliance likelihood score based on the captured at least one steep image and at least one shallow image of the detected vehicle.

26. The system of claim 23, wherein the impairment determination module is further configured to determine an object of interest likelihood score based on the captured at least one steep image and at least one shallow image of the detected vehicle.

27. The system of claim 24, 25 or 26, wherein the impairment determination module is further configured to determine the impairment likelihood score based on the determined mobile device distraction likelihood score and/or the determined seatbelt non-compliance likelihood score and/or the determined object of interest likelihood score.

28. The system of claim 15, wherein the impairment determination module is configured to send the determined impairment likelihood score to a communication device for notifying an enforcement officer.