

Early Warning and Collision Avoidance

This application is entitled to the benefit of the filing dates of United States application nos. 15/994,568, filed on May 31, 2018, 15/994,826, filed on May 31, 2018, 15/994,702, filed on May 31, 2018, 15/994,915, filed on May 31, 2018, 16/222,536 filed on December 17, 2018,
5 15/994,850 filed May 31, 2018 which all claim priority to and the benefit of United States provisional patent application 62/644,725, filed March 19, 2018, the entire contents of which are incorporated here by reference.

Background

This description relates to early warning and collision avoidance.

10 Collision avoidance systems have become abundant. King et al. (US patent publication 2007/0276600 A1, 2007), for example, described placing sensors ahead of an intersection and applying a physics-based decision rule to predict if two vehicles are about to crash at the intersection based on heading and speed.

In Aoude et al. (USA patent 9,129,519 B2, 2015, the entire contents of which are incorporated
15 here by reference) the behavior of drivers is monitored and modeled to allow for the prediction and prevention of a violation in traffic situations at intersections.

Collision avoidance is the main defense against injury and loss of life and property in ground transportation. Providing early warning of dangerous situations aids collision avoidance.

Summary

20 In general, in an example, equipment is located at an intersection of a transportation network. The equipment includes an input to receive data from a sensor oriented to monitor ground transportation entities at or near the intersection. A wireless communication device sends to a device of one of the ground transportation entities, a warning about a dangerous situation at or near the intersection, there is a processor and a storage for instructions executable by the
25 processor to perform actions including the following. A machine learning model is stored that can predict behavior of ground transportation entities at or near the intersection at a current time. The machine learning model is based on training data about previous motion and related behavior of ground transportation entities at or near the intersection. Current motion data

received from the sensor about ground transportation entities at or near the intersection is applied to the machine learning model to predict imminent behaviors of the ground transportation entities. An imminent dangerous situation for one or more of the ground transportation entities at or near the intersection is inferred from the predicted imminent behaviors. The wireless communication device sends the warning about the dangerous situation to the device of one of the ground transportation entities.

Implementations may include one or a combination of two or more of the following features. The wireless communication device sends the warning about the dangerous situation to a sign or other infrastructure presentation device. The warning includes an instruction or command capable of controlling a specific ground transportation entity. The equipment includes a roadside equipment. There is a housing for the equipment and the sensor is attached to the housing. The warning is sent by broadcasting the warning for receipt by any of the ground transportation entities at or near the intersection. The machine learning model includes an artificial intelligence model. The training data and the motion data include at least one of speed, location, or heading. The training data and motion data may also include intent, posture, direction of look, or interaction with other vulnerable road users, such as in a group. The processor is configured to be capable of generating the machine learning model at the equipment. The training data is stored at the equipment. The intersection includes a non-signalized intersection. The intersection includes a signalized intersection. The transportation network includes a road network. The ground transportation entities include vulnerable road users. The ground transportation entities include vehicles. The imminent dangerous situation includes a collision or a near miss. The ground transportation entities include pedestrians crossing a road at a crosswalk. There is another communication device to communicate with a central server. The device of one of the ground transportation entity includes a mobile communication device.

In general, in an example, equipment is located at an intersection of the transportation network. The equipment includes an input to receive data from a sensor oriented to monitor ground transportation entities at or near the intersection. A wireless communication device sends to a device of one of the ground transportation entities, a warning about a dangerous situation at or near the intersection. There is a processor and a storage for instructions executable by the processor to store a machine learning model that can predict behavior of ground transportation

entities at or near the intersection at a current time. The machine learning model is based on training data about previous motion and related behavior of ground transportation entities at or near the intersection. Current motion data received from the sensor about ground transportation entities at or near the intersection is applied to the machine learning model to predict imminent behaviors of the ground transportation entities, including a ground transportation entity the device of which cannot receive a warning from the wireless communication device. An imminent dangerous situation for a ground transportation entity the device of which can receive a warning from the wireless communication device is inferred. The imminent dangerous situation is the result of predicted imminent behavior of the ground transportation entity that cannot receive the warning. The warning about the dangerous situation is sent to the device of the ground transportation entity that can receive the warning from the wireless communication device.

Implementations may include one or a combination of two or more of the following features. The equipment includes a roadside equipment. There is a housing for the equipment and the sensor is attached to the housing. The warning is sent by broadcasting the warning for receipt by any of the ground transportation entities at or near the intersection that can receive the warning. The machine learning model includes an artificial intelligence model. The intersection includes a non-signalized intersection. The intersection includes a signalized intersection. The transportation network includes a road network. The ground transportation entities include vulnerable road users. The ground transportation entities include vehicles. The imminent dangerous situation includes a collision. The ground transportation entity the device of which cannot receive the warning from the wireless communication device includes a vehicle. The ground transportation entity the device of which can receive the warning from the wireless communication device includes a pedestrian crossing a road at a crosswalk. There is another communication device to communicate with a central server. The device of one of the ground transportation entity includes a mobile communication device.

In general, in an example, on board a road vehicle traveling in a ground transportation network messages and data are received including messages from external sources about location, motion, and state of other ground transportation entities, data from on board sensors about road and driving conditions and about the locations of static objects and moving ground transportation entities in the vehicle surroundings, data about quality of driving by a driver of the road vehicle,

and basic safety messages from other ground transportation entities and personal safety messages from vulnerable road users. The received data and messages are fused and applied to an artificial intelligence model to predict an action of a driver of the road vehicle or of a vulnerable road user or a collision risk for the road vehicle or both.

- 5 Implementations may include one or a combination of two or more of the following features. The road vehicle creates a map of the static objects and moving ground transportation entities in the vicinity of the road vehicle. The driver of the road vehicle is alerted to a collision risk. The collision risk is determined based on probabilities of predicted trajectories of nearby other moving ground transportation entities. Basic safety messages and personal safety messages are
10 filtered to reduce the number of alerts provided to the driver of the road vehicle.

In general, in a first aspect, the present disclosure provides a method comprising using electronic sensors located in a vicinity of a crosswalk that crosses a road to monitor an area in and nearby the crosswalk, the electronic sensors generating motion data about vulnerable roadway users who are in or nearby the crosswalk, the motion data including location and direction of the vulnerable
15 roadway users, applying the generated motion data to a machine learning model running in equipment located in the vicinity of the crosswalk to make a prediction that one of the vulnerable roadway users is about to enter the road in or nearby the crosswalk, the prediction being made prior to the vulnerable roadway user entering the road in or nearby the crosswalk, the machine learning model having been trained using motion data generated in the vicinity of the crosswalk
20 and indicative of intent or behavior of vulnerable roadway users who previously had been in or nearby the crosswalk, the motion data comprising location, speed, acceleration, and orientation, and before the vulnerable roadway users has entered the road in or nearby the crosswalk, transmitting a warning to at least one of: a device associated with the vulnerable roadway user, a device associated with another ground transportation entity that is approaching the crosswalk on
25 the road, or a road sign configured to alert the vulnerable roadway user or a driver.

Implementations may include one or a combination of two or more of the following features. The equipment includes a roadside equipment. The vulnerable roadway user includes a pedestrian, animal, or cyclist. The device associated with the vulnerable roadway user includes a smart watch or other wearable, a smart phone, or another mobile device. The other ground
30 transportation entity includes a motorized vehicle. The device associated with the other ground

transportation entity includes a smart phone or another mobile device. The machine learning model is provided to the equipment located in the vicinity of the crosswalk by a remote server through the Internet. The machine learning model is generated at the equipment located in the vicinity of the crosswalk. The machine learning model is trained using motion data generated by
5 the sensors located in the vicinity of the crosswalk. Motion data generated by the sensors located in the vicinity of the crosswalk is sent to a server for use in training the machine learning model.

The motion data generated by the sensors located in the vicinity of the crosswalk is segmented based on corresponding zones in the vicinity of the crosswalk. The electronic sensors are used to generate motion related data representing physical properties of the vulnerable road user.

Trajectory information about the vulnerable road user is derived from motion data generated by the sensor.

In general, in an example, electronic sensors located in a vicinity of an intersection of a ground transportation network are used to monitor the intersection and approaches to the intersection.

The electronic sensors generate motion data about ground transportation entities moving on the approaches or in the intersections. One or more of the ground transportation entities are not

capable of sending basic safety messages to other ground transportation entities in the vicinity of the intersection. Based on the motion data generated by the electronic sensors, virtual basic safety messages are sent to one or more of the ground transportation entities that are capable of receiving the messages. The virtual basic safety messages incorporate information about one or more of the ground transportation entities that are not capable of sending basic safety messages.

The incorporated information in each of the virtual basic safety messages includes at least one of the location, heading, speed, and predicted future trajectory of one of the ground transportation entities that are not capable of sending basic safety messages.

Implementations may include one or a combination of two or more of the following features. The equipment includes a roadside equipment. The incorporated information includes a subset of

information that would be incorporated in a basic safety message generated by the ground transportation entity if it were capable of sending basic safety messages. The generated motion

data is applied to a machine learning model running in equipment located in the vicinity of the intersection to predict a trajectory of the ground transportation entity that is not capable of

sending basic safety messages. One of the ground transportation entities includes a motorized

vehicle. The machine learning model is provided to the equipment located in the vicinity of the intersection by a remote server through the Internet. The machine learning model is generated at

the equipment located in the vicinity of the intersection. The machine learning model is trained

using motion data generated by the sensors located in the vicinity of the intersection. Motion data generated by the sensors located in the vicinity of the intersection is sent to a server for use in

training the machine learning model.

In general, in an example, electronic sensors located in a vicinity of an intersection of a ground transportation network are used to monitor the intersection and approaches to the intersection. The electronic sensors generate motion data about ground transportation entities moving on the approaches or in the intersections. Distinct virtual zones are defined in the intersection and the approaches to the intersection. The motion data is segmented according to corresponding virtual zones to which the generated motion data relates. The generated motion data is applied for each of the respective segments to a machine learning model running in equipment located in the vicinity of the intersection to predict an imminent dangerous situation in the intersection or one of the approaches involving one or more of the ground transportation entities. Before the imminent dangerous situation becomes an actual dangerous situation, a warning is wirelessly transmitted to a device associated with at least one of the involved ground transportation entities.

Implementations may include one or a combination of two or more of the following features. The equipment includes a roadside equipment. The device associated with each of the ground transportation entities includes a wearable, a smart phone, or another mobile device. One of the ground transportation entities includes a motorized vehicle. The machine learning model is provided to the equipment located in the vicinity of the intersection by a remote server through the Internet. The machine learning model is generated at the equipment located in the vicinity of the intersection. The machine learning model is trained using motion data generated by the sensors located in the vicinity of the intersection. The motion data generated by the sensors located in the vicinity of the intersection is sent to a server for use in training the machine learning model. The electronic sensors are used to monitor an area in or nearby a crosswalk that crosses one of the approaches to the intersection. The electronic sensors are used to generate motion related data representing physical properties of a vulnerable road user in the vicinity of the crosswalk. Trajectory information about the vulnerable road user is derived from motion data generated by the sensor. There is a machine learning model for each of the approaches to the intersection. A determination is made whether to transmit the warning based also on motion data generated by sensors with respect to another nearby intersection. A determination is made whether to transmit the warning based also on information received from ground transportation entities moving on the approaches or in the intersection. The intersection is signalized and information about the state of the signals is received. The intersection is not signalized and is

controlled by one or more signs. The defined virtual zones include one or more approaches controlled by the signs. The signs include a stop sign or a yield sign. One of the ground transportation entities includes a rail vehicle.

In general, in an example, equipment is located in or on a ground transportation entity. The equipment includes an input to receive data from sensors in or on the ground transportation entity and oriented to monitor nearby features of a ground transportation network and other information about a context in which the ground transportation entity is traversing the ground transportation network. A wireless communication device receives information about the context. A signal processor applies signal processing to data from the sensor and other information about the context. There is a processor and a storage for instructions executable by the processor to perform actions that include the following: store a machine learning model that can predict behavior of an operator of the ground transportation entity and intent and movement of other ground transportation entities in the vicinity, and apply the current received data from the sensors and other information about the context to predict behavior of the operator and the intent and movement of other ground transportation entities in the vicinity.

Implementations may include one or a combination of two or more of the following features. The equipment includes a roadside equipment. The instructions are executable by the processor to monitor users or occupants of the ground transportation entity. The other information about the context includes emergency broadcasts, traffic and safety messages road side equipment, and messages about safety, locations, and other motion information from other ground transportation entities. The sensors include cameras, range sensors, vibration sensors, microphones, seating sensors, hydrocarbon sensors, sensors of volatile organic compounds and other toxic materials, and kinematic sensors or combinations of them. The instructions are executable by the processor to filter received alerts that the vehicle receives by applying the alerts to a machine learning model to predict which alerts are important in the current location, environmental conditions, driver behavior, vehicle health and status, and kinematics.

In general, in an example, motion data are acquired for unconnected ground transportation entities moving in a transportation network. Virtual safety messages incorporating information about the motion data for the unconnected ground transportation entities at sent to connected ground transportation entities in the vicinity of the unconnected ground transportation entities.

Implementations may include one or a combination of two or more of the following features. the virtual safety messages are substitutes for safety messages that would be sent by the unconnected ground transportation entities if they were connected. The unconnected ground transportation entities include vehicles and the virtual safety messages are substitutes for basic safety messages.

- 5 The unconnected ground transportation entities include vulnerable road users and the virtual safety messages are substitutes for personal safety messages. The motion data are detected by infrastructure sensors.

In general, in an example, equipment located at an intersection of a transportation network includes inputs to receive data from sensors oriented to monitor ground transportation entities at
10 or near the intersection. The data from each of the sensors represents at least one location or motion parameter of at least one of the ground transportation entities. The data from each of the sensors is expressed in a native format. The data received from at least two of the sensors is inconsistent with respect to the location or motion parameters or the native formats or both. There is a storage for instructions executable by a processor to convert the data from each of the
15 sensors into data having a common format independent of the native formats of the data of the sensors. The data having the common format is incorporated into a global unified representation of the ground transportation entities being monitored at or near the intersection. The global unified representation includes the location, speed, and heading of each of the ground transportation entities. Relationships of locations and motions of two of the ground
20 transportation entities are determined using the global unified representation. A dangerous situation is predicted involving the two ground transportation entities, and a message is sent to at least one of the two ground transportation entities alerting it to the dangerous situation.

Implementations may include one or a combination of two or more of the following features. The sensors include at least two of: radar, lidar, and a camera. The data received from one of the
25 sensors includes image data of a field of view at successive moments. The data received from one of the sensors includes points of reflection in 3D space. The data received from one of the sensors includes distance from the sensor and speed. The global unified representation represents locations of the ground transportation entities in a common reference frame. Two sensors from which the data is received are mounted in fixed positions at or near the intersection and have at
30 least partially non-overlapping fields of view. One of the sensors includes radar and the

converting of the data includes determining locations of ground transportation entities from a known location of the radar and distances from the radar to the ground transportation entities. One of the sensors includes a camera and the converting of the data includes determining locations of ground transportation entities from a known location, direction of view, and tilt of the camera and the locations of the ground transportation entities within an image frame of the camera.

In general, in a second aspect, the present disclosure provides an apparatus comprising equipment located at a level crossing of a transportation network, the level crossing comprising an intersection of a road, a pedestrian crossing, and a rail line, the equipment comprising inputs to receive data from sensors oriented to monitor road vehicles and pedestrians at and near the level crossing and to receive phase and timing data for signals on the road and on the rail line, a wireless communication device to send to a device of one of the ground transportation entities, pedestrians, or rail vehicles on the rail line, a warning about a dangerous situation at or near the level crossing, a processor, and a storage for instructions executable by the processor to store a machine learning model that can predict behavior of ground transportation entities at or near the level crossing at a current time, the machine learning model being based on training data about previous motion and related behavior of road vehicles and pedestrians at or near the intersection, and previous phase and timing data for signals on the road and on the rail line, apply current motion data received from the sensors about road vehicles and pedestrians at or near the level crossing and current phase and timing data for signals on the road and on the rail line to the machine learning model to predict imminent behaviors of the road vehicles and pedestrians, infer an imminent dangerous situation for a rail vehicle on the rail line at or near the intersection from the predicted imminent behaviors, cause the wireless communication device to send the warning about the dangerous situation to a device of at least one of the road vehicles, pedestrians, and rail vehicle.

Implementations may include one or a combination of two or more of the following features. The warning is sent to an on-board equipment of the rail vehicle. The rail line is on a segregated right of way. The rail line is not on a segregated right of way. The equipment includes a roadside equipment. The warning is sent by broadcasting the warning for receipt by any of the ground

transportation entities, pedestrians, or rail vehicles at or near the level crossing. The imminent dangerous situation includes a collision or a near miss.

In general, in a third aspect, the present disclosure provides a method comprising receiving data from infrastructure sensors representing positions and motions of road vehicles being driven or
5 pedestrians walking in a ground transportation network,

receiving data in basic safety messages or virtual basic safety messages and personal safety message and virtual personal safety messages about states of the road vehicles and pedestrians, the basic safety messages, virtual basic safety messages, personal safety messages, and virtual personal safety messages comprising location, heading, and speed information about the road vehicles or pedestrians, the basic safety messages and the personal safety messages being
5 received from the road vehicles or pedestrians, the virtual basic safety messages and the virtual personal safety messages comprising location, heading, and speed information about the road vehicles and the pedestrians reconstructed from the data received from the infrastructure sensors, applying the received data from the infrastructure sensors and the data from the basic safety
10 messages, the virtual basic safety messages, the personal safety messages, and the virtual personal safety messages to a machine learning model trained to identify dangerous driving or walking behavior of one of the road vehicles or pedestrians, and automatically reporting the dangerous driving or walking behavior to authorities.

Implementations may include one or a combination of two or more of the following features. The
15 road vehicles are identified based on plate number recognition. The pedestrians are identified based on biometric recognition. The road vehicles or pedestrians are identified based on social networking.

These and other aspects, features, and implementations can be expressed as methods, apparatus, systems, components, program products, methods of doing business, means or steps for
20 performing a function, and in other ways.

These and other aspects, features, and implementations will become apparent from the following descriptions, including the claims.

Description

Figures 1, 2, 3, and 15 are block diagrams.

25 Figures 4, 5, 8 through 11, 13, 14, 17, and 18 are schematic views of road networks from above.

Figures 6 and 7 are annotated perspective views of intersections.

Figures 12 and 16 are schematic side and perspective views of road networks.

With advancements in sensor technologies and computers, it has become feasible to predict (and to provide early warning of) dangerous situations and in that way to prevent collisions and near misses of ground transportation entities (that is, to enable collision avoidance) in the conduct of ground transportation.

- 5 We use the term “ground transportation” broadly to include, for example, any mode or medium of moving from place to place that entails contact with the land or water on the surface of the earth, such as walking or running (or engaging in other pedestrian activities), non-motorized vehicles, motorized vehicles (autonomous, semi-autonomous, and non-autonomous), and rail vehicles.

We use the term “ground transportation entity” (or sometimes simply “entity”) broadly to include, for example, a person or a discrete motorized or non-motorized vehicle engaged in a mode of ground transportation, such as a pedestrian, bicycle rider, boat, car, truck, tram, streetcar, or train, among others. Sometimes we use the terms “vehicle” or “road user” as shorthand references to a ground transportation entity.

We use the term “dangerous situation” broadly to include, for example, any event, occurrence, sequence, context, or other situation that may lead to imminent property damage or personal injury or death and that may be reducible or avoidable. We sometimes use the term “hazard” interchangeably with “dangerous situation.” We sometimes use the word “violation” or “violate” with respect to behavior of an entity that has, may, or will lead to a dangerous situation.

In some implementations of the technology that we discuss here a ground transportation network is being used by a mix of ground transportation entities that do not have or are not using transportation connectivity and ground transportation entities that do have and are using transportation connectivity.

We use the term “connectivity” broadly to include, for example, any capability a ground transportation entity to (a) be aware of and act on knowledge of its surroundings, other ground transportation entities in its vicinity, and traffic situations relevant to it, (b) broadcast or otherwise transmit data about its state, or (c) both (a) and (b). The data transmitted can include its location, heading, speed, or internal states of its components relevant to a traffic situation. In some cases, the awareness of the ground transportation entity is based on wirelessly received data about other ground transportation entities or traffic situations relevant to the operation of the ground transportation entity. The received data can originate from the other ground transportation entities or from infrastructure devices, or both. Typically connectivity involves sending or receiving data in real time or essentially real time or in time for one or more of the ground transportation entities to act on the data in a traffic situation.

We use the term “traffic situation” broadly to include any circumstance in which two or more ground transportation entities are operating in the vicinity of one another and in which the operation or status of each of the entities can affect or be relevant to the operation or status of the others.

We sometimes refer to a ground transportation entity that does not have or is not using connectivity or aspects of connectivity as a “non-connected ground transportation entity” or simply a “non-connected entity.” We sometimes refer to a ground transportation entity that has and is using connectivity or aspects of connectivity as a “connected ground transportation entity”
5 or simply a “connected entity.”

We sometimes use the term “cooperative entity” to refer to a ground transportation entity that broadcasts data to its surroundings including location, heading, speed, or states of on board safety systems (such brakes, lights, and wipers), for example.

We sometimes use the term “non-cooperative entity” to refer to a ground transportation entity
10 that does not broadcast to its surroundings one or more types of data, such as its location, speed, heading, or state.

We sometimes use the term “vicinity” of a ground transportation entity broadly to include, for example, an area in which a broadcast by the entity can be received by other ground transportation entities or infrastructure devices. In some cases, the vicinity varies with location of
15 the entity and the number and characteristics of obstacles around the entity. An entity traveling on an open road in a desert will have a very wide vicinity since there are no obstacles to prevent a broadcast signal from the entity from reaching long distances. Conversely, the vicinity in an urban canyon will be diminished by the buildings around the entity. Additionally, there may be sources of electromagnetic noise that degrade the quality of the broadcast signal and therefore
20 the distance of reception (the vicinity).

As shown in figure 14, the vicinity of an entity 7001 traveling along a road 7005 can be represented by concentric circles with the outermost circle 7002 representing the outermost extent of the vicinity. Any other entity that lies within the circle 7002 is in the vicinity of entity 7001. Any other entity that lies outside the circle 7002 is outside the vicinity of, and unable to
25 receive a broadcast by, the entity 7001. The entity 7001 would be invisible to all entities and infrastructure devices outside its vicinity.

Typically, cooperative entities are continuously broadcasting their state data. Connected entities in the vicinity of a broadcasting entity are able to receive these broadcasts and can process and act on the received data. If, for example, a vulnerable road user has a wearable device that can

receive broadcasts from an entity, say an approaching truck, the wearable device can process the received data and let the vulnerable user know when it is safe to cross the road. This operation occurs without regard to the locations of the cooperative entity or the vulnerable user relative to a “smart” intersection as long as the user’s device can receive the broadcast, i.e., is within the vicinity of the cooperative entity.

We use the term “vulnerable road users” or “vulnerable road users” broadly to include, for example, any user of roadways or other features of the road network who is not using a motorized vehicle. vulnerable road users are generally unprotected against injury or death or property damage if they collide with a motorized vehicle. In some examples, vulnerable road users could be people walking, running, cycling or performing any type of activity that puts them at risk of direct physical contact by vehicles or other ground transportation entities in case of a collisions.

In some implementations, the collision avoidance technologies and systems described in this document (which we sometimes refer to simply as the “system”) use sensors mounted on infrastructure fixtures to monitor, track, detect, and predict motion (such as speed, heading, and position), behavior (e.g., high speed), and intent (e.g., will violate the stop sign) of ground transportation entities and drivers and operators of them. The information provided by the sensors (“sensor data”) enables the system to predict dangerous situations and provide early warning to the entities to increase the chances of collision avoidance.

We use the term “collision avoidance” broadly to include, for example, any circumstance in which a collision or a near miss between two or more ground transportation entities or between a ground transportation entity and another object in the environment that may result from a dangerous situation, is prevented or in which chances of such an interaction are reduced.

We use the term “early warning” broadly to include, for example, any notice, alert, instruction, command, broadcast, transmission, or other sending or receiving of information that identifies, suggests, or is in any way indicative of a dangerous situation and that is useful for collision avoidance.

Road intersections are prime locations where dangerous situations can happen. The technology that we describe here can equip intersections with infrastructure devices including sensors,

computing hardware and intelligence to enable simultaneous monitoring, detection, and prediction of dangerous situations. The data from these sensors is normalized to a single frame of reference and then is processed. Artificial intelligence models of traffic flow along different approaches to the intersection are constructed. These models help, for example, entities that are
5 more likely to violate traffic rules. The models are set up to detect the dangerous situations before the actual violations and therefore can be considered as predictions. Based on a prediction of a dangerous situation, an alert is sent from the infrastructure devices at the intersection to all connected entities in the vicinity of the intersection. Every entity that receives an alert, processes the data in the alert and performs alert filtering. Alert filtering is a process of discarding or
10 disregarding alerts that are not beneficial to the entity. If an alert is considered beneficial (i.e., is not disregarded as a result of the filtering), such as an alert of an impending collision, the entity either automatically reacts to the alert (such as by applying brakes), or a notification is presented to the driver or both.

The system can be used on, but is not limited to, roadways, waterways, and railways. We
15 sometimes refer to these and other similar transportation contexts as “ground transportation networks.”

Although we often discuss the system in the context of intersections, it can also be applied to other contexts.

We use the term “intersection” broadly to include, for example, any real-world arrangement of
20 roads, rails, water bodies, or other travel paths for which two or more ground transportation entities traveling along paths of a ground transportation network could at some time and location occupy the same position producing a collision.

The ground transportation entities using a ground transportation network move with a variety of speeds and may reach a given intersection at different speeds and times of the day. If the speed
25 and distance of an entity from the intersection is known, dividing the distance by the speed (both expressed in the same unit system) will give the time of arrival at the intersection. However, since the speed of will change due, for example, to traffic conditions, speed limits on the route, traffic signals, and other factors, the expected time of arrival at the intersection changes

continuously. This dynamic change in expected time of arrival makes it impossible to predict the actual time of arrival with 100% confidence.

To account for the factors affecting the motion of an entity requires applying a large number of relationships between the speed of the entity and the various affecting factors. The absolute
5 values of the state of motion of an entity can be observed by a sensor tracking that entity either from the entity or from an external location. The data captured by these sensors can be used to model the patterns of motion, behaviors, and intentions of the entities. Machine learning can be used to generate complex models from vast amounts of data. Patterns that cannot be modeled using kinematics of the entities directly can be captured using machine learning. A trained model
10 can predict whether an entity is going to move or stop at a particular point by using that entity's tracking data from the sensors tracking them.

In other words, in addition to detecting information about ground transportation entities directly from the sensor data, the system uses artificial intelligence and machine learning to process vast amounts of sensor data to learn the patterns of motion, behaviors, and intentions of ground
15 transportation entities, for example, at intersections of ground transportation networks, on approaches to such intersections, and at crosswalks of ground transportation networks. Based on the direct use of current sensor data and on the results of applying the artificial intelligence and machine learning to the current sensor data, the system produces early warnings such as alerts of dangerous situations and therefore aids collision avoidance. With respect to early warnings in the
20 form of instructions or commands, the command or instruction could be directed to a specific autonomous or human-driven entity to control the vehicle directly. For example, the instruction or command could slow down or stop an entity being driven by a malevolent person who has been determined to be about to run a red light for the purpose of trying to hurt people.

The system can be tailored to make predictions for that particular intersection and to send alerts
25 to the entities in the vicinity of the device broadcasting the alerts. For this purpose, the system will use sensors to derive data about the dangerous entity and pass the current readings from the sensors through the trained model. The output of the model then can predict a dangerous situation and broadcast a corresponding alert. The alert, received by connected entities in the vicinity, contains information about the dangerous entity so that the receiving entity can analyze
30 that information to assess the threat posed to it by the dangerous entity. If there is a threat, the

receiving entity can either take action itself (e.g., slowing down) or notify the driver of the receiving entity using a human machine interface based on visual, audio, haptic, or any kind of sensory stimulation. An autonomous entity may take action itself to avoid a dangerous situation.

The alert can also be sent directly through the cellular or other network to a mobile phone or
5 other device equipped to receive alerts and possessed by a pedestrian. The system identifies potential dangerous entities at the intersection and broadcasts (or directly sends) alerts to a pedestrian's personal device having a communication unit. The alert may, for example, prevent a pedestrian from entering a crosswalk and thus avoid a potential accident.

The system can also track pedestrians and broadcast information related to their state (position,
10 speed, and other parameters) to the other entities so that the other entities can take action to avoid dangerous situations.

As shown in figure 1, the system includes at least the following types of components:

1. Roadside Equipment (RSE) 10 that includes or makes use of sensors 12 to monitor, track, detect, and predict motion (such as speed, heading, and position), behavior (e.g., high speed), and
15 intent (e.g., will violate the stop sign) of ground transportation entities 14. The RSE also includes or can make use of a data processing unit 11 and data storage 18. The ground transportation entities exhibit a wide range of behavior which depends on the infrastructure of the ground transportation network as well as the states of the entities themselves, the states of the drivers, and the states of other ground transportation entities. To capture the behaviors of the entities the
20 RSE collects information from the sensors, other RSEs, OBEs, OPEs, local or central servers, and other data processing units. The RSE also saves the data received by it as well as may save the processed data at some or all the steps in the pipeline.

The RSE may save the data on a local storage device or a remote storage. The collected data is processed in real time using predefined logic or logic based on the data collected dynamically
25 which means that the RSE can update its own logic automatically. The data can be processed over a single processing unit or a cluster of processing units to get results faster. The data can be processed on a local or remote processing unit or a local or remote cluster of processing units. The RSE can use a simple logic or a sophisticated model trained on the collected data. The model can be trained locally or remotely.

The RSE may preprocess data before using the trained model to filter outliers. The outliers can be present due to noise in the sensor, reflections or due to some other artifact. The resulting outliers can lead to false alarms which can affect the performance of the whole RSE. The filtration methods can be based on the data collected by the RSE, OBEs, OPEs, or online
5 resources. The RSE may interface with other controllers such as traffic light controllers at the intersection or other location to extract information for use in the data processing pipeline.

The RSE also includes or can make use of communication equipment 20 to communicate by wire or wireless with other RSEs, and with OBEs, OPEs, local or central servers, and other data processing units. The RSE can use any available standard for communication with other
10 equipment. The RSE may use wired or wireless Internet connections for downloading and uploading data to other equipment, the cellular network to send and receive messages from other cellular devices, and a dedicated radio device to communicate to infrastructure devices and other RSEs at the intersection or other location.

An RSE can be installed next to different kinds of intersections. For example, at a signalized
15 intersection (e.g., an intersection in which traffic is controlled by a light), an RSE 10 is installed near the traffic light controllers 26 either in the same enclosure or within a nearby enclosure. Data (such as traffic light phase and timing) is meant to flow 28 between the traffic light controllers and the RSE. At a non-signalized intersection, the RSE 10 is usually located to make it easy to connect it to the sensors 12 that are used to monitor the roads or other features of the
20 ground transportation network in the vicinity of the intersection. The proximity of RSE with the intersection helps in maintaining a low latency system which is crucial for providing maximum time to the receiving ground units to respond to an alert.

2. Onboard Equipment (OBE) 36 mounted on or carried by or in the ground transportation entities 14, which includes sensors 38 that determine location and kinematics (motion data) of
25 the entities in addition to safety related data about the entities. OBEs also include data processing units 40, data storage 42, and communication equipment 44 that can communicate wirelessly with other OBEs, OPEs, RSEs, and possibly servers and computing units.

3. On Person Equipment (OPE) 46 which can be, but is not limited to, a mobile phone, wearable device, or any other device that is capable of being worn by, held by, attached to, or otherwise

interfacing with a person or animal. OPEs can include or be coupled to data processing units 48, data storage 50, and communication equipment 52 if needed. In some implementations, an OPE serves as a dedicated communication unit for a non-vehicular vulnerable road user. In some cases, the OPE can also be used for other purposes. The OPE may have a component to provide
5 visual, audio, or haptic alerts to the vulnerable road user.

Vulnerable road user can include pedestrians, cyclists, road workers, people on wheelchairs, scooters, self-balancing devices, battery powered personal transporters, animal driven carriages, guide or police animals, farm animals, herds, and pets.

Typically an OPE is in the possession of the vulnerable road user and is capable of sending and
10 receiving messages. An OPE can be attached to or integrated with a mobile phone, tablet, personal transporter, bicycle, wearable device (watch, bracelet, anklet, for example), or attached to a pet collar.

Messages sent by an OPE can include kinematic information associated with the vulnerable road user including, but not limited to, time of day, 3D position, heading, velocity, and acceleration.
15 Sent messages can also carry data representing the alertness level, current behavior, and future intents of the vulnerable road user, e.g. that the vulnerable road user is currently crossing the crosswalk, is listening to music, or is going to cross the crosswalk. Among other things, the message may convey the blob size or data size of the vulnerable road user, whether there are external devices with the vulnerable road user (e.g., a stroller, a cart, or other device), whether
20 the vulnerable road user has a disability or is using any personal assistance. The message may convey the category of worker if the vulnerable road user is a worker and may also describe the type of activity being done by the worker. When a cluster of similar vulnerable road users (say, a group of pedestrians) have similar characteristics, a single message can be sent to avoid multiple message broadcasts.

25 Typically, the messages received by an OPE are alert messages from a roadside equipment or from an entity. The OPE can act on the received messages by alerting the vulnerable road user. The alert message will carry data useful in providing a custom alert for the vulnerable road user. For example, the alert to the vulnerable road user may showcase a type of dangerous situation

and suggest possible actions. The OPE can apply alert filtering to all received messages and present only relevant messages to the vulnerable road user.

Alert filtering is based on the outcome of applying a learning algorithm to historical data associated with the OPE which enables custom-tailoring the alert filtering to each vulnerable road user. The OPE learning algorithm tracks the responses of the vulnerable road user to received alerts and tailors future alerts to attain the best response time and the best attention from vulnerable road user. The learning algorithm can also be applied to data carried in sent messages.

4. Data storage servers 54 which can be but are not limited to cloud storage, local storage, or any other storage facility that allows for storage and retrieval of data. The data storage servers are accessible by RSEs, computing units, and potentially by OBEs, OPEs, and data servers, for the purpose of storing data related to early warning and collision avoidance, for example. The data storage servers are accessible from RSEs and potentially from OBEs, OPEs, and data servers, for the purpose of fetching stored data. The data can be raw sensor data, processed data by a processing unit or any other information generated by the RSEs, OBEs and OPEs.

15 Sensors at an intersection, which monitor ground transportation entities continuously, can generate a large amount of data every day. The volume of this data depends on the number and types of the sensors. The data is both processed in real time and saved for future analysis requiring data storage units (e.g., hard disk drives, solid state drives, and other mass storage devices) locally such as at the intersection. The local storage devices will get filled up in a period depending on their storage capacity, the volume of generated data, and the rate at which it is generated. To preserve the data for future use, the data is uploaded to a remote server which has a lot more capacity. The remote server may upgrade the storage capacity on demand as needed. The remote server may use a data storage device similar to the local storage (e.g., a hard disk drive, a solid state drive, or other mass storage device) accessible through a network connection.

25 The data stored locally and on the server for future analysis may include the data broadcast by the ground transportation entities and received by the RSE which is saved for future analysis. The stored data can be downloaded from the servers or other remote source for processing on the RSE. For example, the machine learning model of the intersection where the RSE is located may

be stored at the server or in other remote storage, and downloaded by the RSE to use for analyzing the current data received at the RSE from local sources.

5 Computing units 56 which are powerful computing machines located in the cloud or locally (for example as part of an RSE) or a combination of those. Among other functions, the computing units process the available data to generate predictions, machine learning based models of motions, behaviors, and intents of the vehicles, pedestrians, or other ground transportation entities using the transportation network. Each of the computing units can have dedicated hardware to process corresponding types of data (e.g., a graphics processing unit for processing images). In case of heavy processing loads, the computing unit in the RSE may become overloaded. This may happen, for example, when additional data generation units (e.g. sensors) are added to the system producing a computational overload. The overload can also occur if the logic running in the computing unit is replaced with more computationally intensive logic. An overload may be caused by an increase in the number of ground transportation entities being tracked. When a local computational overload happens, the RSE can offload some of the tasks to another computing unit. The other computing unit could be nearby the RSE or remote, such as a server. Computational tasks can be prioritized and tasks which are not time critical can be completed at the other computing unit and the results retrieved by the local computing unit. For example, the computing unit in the RSE can request another computing unit to run a job for analyzing saved data and training a model using the data. The trained model will then be downloaded by the computing unit at the RSE to store and use there.

The computing unit at the RSE can use a other small computing units to perform a computationally intensive job more efficiently and saving time. The available computing units are used wisely to perform the most tasks in the least time, for example, by dividing the tasks between the RSE computing units and the other available computing units. A computing unit can also be attached as an external device to an RSE to add more computational capability to the computing unit in the RSE. The externally attached computing unit can have the same or a different architecture as compared to the computing unit in the RSE. The externally attached computing unit may communicate with the existing computing unit using any available communication port. The RSE computing unit can request more computational power from the external computing unit as needed.

The rest of this document will explain in detail the roles and functions of the components above in the system, among other things.

Roadside Equipment (RSE)

As shown in figure 2, an RSE may include, but not be limited to, the following components:

- 5 1. One or more communication units 103, 104 which enable the reception or transmission or both of motion data and other data related to ground transportation entities and traffic safety data, from and to nearby vehicles or other ground transportation entities, infrastructure, and remote servers and data storage systems 130. In some cases, this type of communication is known as infrastructure-to-everything (I2X), which includes but is not limited to infrastructure-to-vehicles
10 (I2V), infrastructure-to-pedestrians (I2P), infrastructure-to-infrastructure (I2I), and infrastructure-to-devices (I2D), and combinations of them. The communication may be wireless or wired and comply with a wide variety of communication protocols.
2. Communication unit 103 is used for communication with ground transportation entities and unit 104 is used for communication through the Internet with remote servers and data storage
15 systems 130.
3. Local storage 106 for storing programs, intersection models, and behavior and traffic models. It may also be used for temporary storage of data collected from the sensors 101.
4. Sensors 101 and sensor controllers 107 which allow for the monitoring of (e.g., generating of
20 data about) moving subjects such as ground transportation entities typically near the RSE. The sensors may include, but are not limited to, cameras, radars, lidars, ultrasonic detectors or any other hardware that can sense or infer from sensed data the distance to, speed, heading, location, or combinations of them, among other things, of a ground transportation entity. Sensor fusion is performed using aggregations or combinations of data from two or more sensors 101.
5. A location receiver (102) (such as a GPS receiver) that provides localization data (e.g.,
25 coordinates of the location of the RSE)) and helps with correcting localization errors in the localization of ground transportation entities.
6. A processing unit 105 that will acquire and use the data generated from the sensors as well as incoming data from the communication units 103, 104. The processing unit will process and

store the data locally and, in some implementations, transmit the data for remote storage and further processing. The processing unit will also generate messages and alerts that are broadcast or otherwise sent through wireless communication facilities to nearby pedestrians, motor vehicles, or other ground transportation entities, and in some cases to signs or other infrastructure presentation devices. The processing unit will also periodically report the health and status of all the RSE systems to a remote server for monitoring.

7. Expansion connector 108 that allows for control and communication between the RSE and other hardware or other components such as temperature and humidity sensors, traffic light controllers, other computing units as described above, and other electronics that may become available in the future.

Onboard Equipment (OBE)

The onboard equipment typically may be original equipment for a ground transportation entity or added to the entity by a third-party supplier. As shown in figure 3, OBE may include, but is not limited to, the following components:

1. A communication unit 203 that enables the sending and receiving, or both, of data to and from nearby vehicles, pedestrians, cyclists, or other ground transportation entities, and infrastructure, and combinations of them. The communication unit also allows for the transmission or reception (or both) of data between the vehicle or other ground transportation entity and a local or remote server 212 for machine learning purposes and for remote monitoring of the ground transportation entity by the server. In some cases, this type of communication is known as vehicle-to-everything (V2X), which includes but is not limited to vehicles-to-vehicles (V2V), vehicles-to-pedestrians (V2P), vehicle-to-infrastructure (V2I), vehicle-to-devices V2D), and combinations of them. The communication may be wireless or wired and comply with a wide variety of communication protocols.

Communication unit 204 will allow the OBE to communicate through the Internet with remote servers for program update, data storage and data processing.

2. Local storage 206 for storing programs, intersection models, and traffic models. It may also be used for temporary storage of data collected from the sensors 201.

3. Sensors 201 and sensor controllers 207 that may include, but are not limited to, external cameras, lidars, radars, ultrasonic sensors or any device that may be used to detect nearby objects or people or other ground transportation entities. Sensors 201 may also include additional kinematic sensors, global positioning receivers, and internal and local microphones and cameras.
- 5 4. A location receiver 202 (such as a GPS receiver) that provides localization data (e.g., coordinates of the location of the ground transportation entity).
5. A processing unit 205 which acquires, uses, generates, and transmits data, including consuming data from and sending data to the communication unit as well as consuming data from sensors in or on the ground transportation entity.
- 10 6. Expansion connectors 208 that allows for control and communication between the OBE and other hardware.
7. An interface unit that can be retrofit or integrated into a head-unit, steering wheel, or driver mobile device in one or more ways such as using visual, audible, or haptic feedback).

Smart OBE (SOBE)

- 15 In a world where all vehicles and other ground transportation entities are connected entities, each vehicle or other ground transportation entity could be a cooperative entity with the others and could report its current location, safety status, intent, and other information to the others. Presently, almost all vehicles are not connected entities, cannot report such information to other ground transportation entities, and are operated by people with different levels of skill,
- 20 wellbeing, stress, and behavior. Without such connectivity and communication, predicting a vehicle's or ground transportation entity's next move becomes difficult and that translates to a diminished ability to implement collision avoidance and to provide early warnings.
- A smart OBE monitors the surroundings and users or occupants of the ground transportation entity. It also keeps tabs on the health and status of the different systems and subsystems of the
- 25 entity. The SOBE monitors the external world by listening to, for example, the radio transmissions from emergency broadcasts, traffic and safety messages from nearby RSE, and messages about safety, locations, and other motion information from other connected vehicles or other ground transportation entities. The SOBE also interfaces with on board sensors that can watch the road and driving conditions such as cameras, range sensors, vibration sensors,

microphones, or any other sensor that allows of such monitoring. A SOBE will also monitor the immediate surroundings and create a map of all the static and moving objects.

A SOBE can also monitor the behavior of the users or occupants of the vehicle or other ground transportation entity. The SOBE uses microphones to monitor the quality of the conversation. It
5 can also use other sensors such as seating sensors, cameras, hydrocarbon sensors, and sensors of volatile organic compounds and other toxic materials. It can also use kinematic sensors to measure the reaction and behavior of the driver and, from that, infer the quality of driving.

SOBE also receives vehicle-to-vehicle messages (e.g., basic safety messages (BSMs)) from other ground transport entities and vehicle-to-pedestrian messages (e.g., personal safety messages
10 (PSMs)) from vulnerable road users.

The SOBE will then fuse the data from this array of sensors, sources, and messages. It will then apply the fused data to an artificial intelligence model that is not only able to predict the next action or reaction of the driver or user of the vehicle or other ground transportation entity or vulnerable road user, but also be able to predict the intent and future trajectories and associated
15 near-miss or collision risks due to other vehicles, ground transportation entities and vulnerable road users nearby. For example, an SOBE can use the BSMs received from a nearby vehicle to predict that the nearby vehicle is about to enter into a lane change maneuver that creates a risk to its own host vehicle, and can alert the driver of an imminent risk. The risk is computed by the SOBE based on the probability of the various future predicted trajectories of the nearby vehicle
20 (e.g., going straight, changing lane to the right, changing lane to the left), and the associated risk of collision with the host vehicle for each of those trajectories. If the risk of collision is higher than a certain threshold, then the warning is displayed to the driver of the host vehicle.

Machine learning is typically required to predict intent and future trajectories due to the complexity of human driver behavior modeling, which is further impacted by external factors
25 (e.g., changing environmental and weather conditions).

A SOBE is characterized by having powerful computational abilities to be able to process the large number of data feeds some of which provide megabytes of data per second. The quantity of data available is also proportional to the level of detail required from each sensor.

A SOBE will also have powerful signal processing equipment to be able to pull useful information from an environment that is known to have high (signal) noise levels and low signal to noise ratios. SOBE will also protect the driver from the massive number of alerts that the vehicle is receiving by providing smart alert filtering. The alert filtering is the result of the machine learning model which will be able to tell which alert is important in the current location, environmental conditions, driver behavior, vehicle health and status, and kinematics.

Smart OBEs are important for collision avoidance and early warning and for having safer transportation networks for all users and not only for the occupants or users of vehicles that include SOBEs. SOBEs can detect and predict the movements of the different entities on the road and therefore aid collision avoidance.

On Person Equipment (OPE)

As mentioned earlier, on person equipment (OPE) includes any device that may be held by, attached to, or otherwise interface directly with a pedestrian, jogger, or other person who is a ground transportation entity or otherwise present on or making use of a ground transportation network. Such a person may be vulnerable road user susceptible to being hit by a vehicle, for example. OPEs may include, but not be limited to, mobile devices (for example, smart phones, tablets, digital assistants), wearables (e.g., eyewear, watches, bracelets, anklets), and implants. Existing components and features of OPEs can be used to track and report location, speed, and heading. An OPE may also be used to receive and process data and display alerts to the user in various modes (visual, sound, haptic, for example).

Honda has developed a communication system and method for V2P applications focused on direct communication between a vehicle and a pedestrian using OPEs. In one case, the vehicle is equipped with an OBE to broadcast a message to a surrounding pedestrian's OPE. The message carries the vehicle's current status including vehicle parameters, speed, and heading, for example. For example, the message could be a basic safety message (BSM). If needed the OPE will present an alert to the pedestrian, tailored to the pedestrian's level of distraction, about a predicted dangerous situation in order to avoid a collision. In another case, the pedestrian's OPE broadcasts a message (such as a personal safety message (PSM)) to a surrounding vehicle's OBE that the pedestrian might cross the vehicle's intended path. If needed, the vehicle's OBE will

display an alert to the vehicle user about a predicted hazard in order to avoid a collision. See Strickland, Richard Dean, et al. "Vehicle to pedestrian communication system and method." U.S. Patent 9,421,909.

5 The system that we describe here, uses an I2P or I2V approach using sensors external to the vehicle and the pedestrian (mainly on infrastructure) to track and collect data on pedestrians and other vulnerable road users. For example the sensors can track pedestrians crossing a street and vehicles operating at or near the crossing place. The data collected will in turn be used to build predictive models of pedestrian and vehicle driver intents and behaviors on roads using rule-based and machine learning methods. These models will help analyze the data collected and
10 make predictions of pedestrian and vehicle paths and intents. If a hazard is predicted, a message will be broadcast from the RSE to the OBE or the OPE or both, alerting each entity of the intended path of the other and allowing each of them to take a pre-emptive action with enough time to avoid the collision.

Remote Computing (Cloud Computing and Storage)

15 The data collected from the sensors connected to or incorporated in the RSEs, the OBEs, and the OPEs needs to be processed so that effective mathematical machine learning models can be generated. This processing requires a lot of data processing power to reduce the time needed to generate each model. The required processing power is much more than what is typically available locally on the RSE. To address this, the data can be transmitted to a remote computing
20 facility that provides the power needed and can scale on demand. We refer to the remote computing facility as a "remote server" which aligns with the nomenclature used in computing literature. In some cases, it may be possible to perform part or all of the processing at the RCEs by equipping them with high-powered computing capabilities.

Rule Based Processing

25 Unlike artificial intelligence and machine learning techniques, rule-based processing can be applied at any time without the need for data collection, training, and model building. Rule-based processing can be deployed from the beginning of operation of the system, and that it what is typically done, until enough training data has been acquired to create machine learning models. After a new installation, rules are setup to process incoming sensor data. This is not only useful

to improve road safety but also is a good test case to make sure that all the components of the system are working as expected. Rule based processing can be also added and used later as an additional layer to capture rare cases for which machine learning might not be able to make accurate predictions. Rule-based approaches are based on simple relationships between collected data parameters (e.g., speed, range, and others). Rule-based approaches could also provide a baseline for the assessment of the performance of machine learning algorithms.

In rule-based processing, a vehicle or other ground transportation entity traversing part of a ground transportation network is monitored by sensors. If its current speed and acceleration exceed a threshold that would prevent it from stopping before a stop bar (line) on a road, for example, an alert is generated. A variable region is assigned to every vehicle or other ground transportation entity. The region is labeled as a dilemma zone in which the vehicle has not been yet labeled as a violating vehicle. If the vehicle crosses the dilemma zone into the danger zone because its speed or acceleration or both exceed predefined thresholds, the vehicle is labeled as a violating entity and an alert is generated. The thresholds for speed and acceleration are based on physics and kinematics and vary with each ground transportation entity that approaches the intersection, for example.

Two traditional rule-based approaches are 1) static TTI (Time-To-Intersection), and 2) static RDP (Required Deceleration Parameter). See Aoude, Georges S., et al. "Driver behavior classification at intersections and validation on large naturalistic data set." IEEE Transactions on Intelligent Transportation Systems 13.2 (2012): 724-736.

Static TTI (Time-To-Intersection) uses the estimated time to arrive at the intersection as the classification criteria. In its simplest form, TTI is computed as $TTI = \frac{r}{v}$, where r is distance to the crossing line at the intersection, and v is the current speed of the vehicle or other ground transportation entity. The vehicle is classified as dangerous if $TTI < TTI_{req}$, where TTI_{req} is the time required for the vehicle to stop safely once braking is initiated. The TTI_{req} parameter reflects the conservativeness level of the rule-based algorithm. The TTI is computed on the onset of braking, identified as when the vehicle deceleration crosses a deceleration threshold (e.g., -0.075g). If a vehicle never crosses this threshold, the classification is performed at a specified last resort time, which typically ranges from 1s to 2s of estimated remaining time to arrive at the intersection.

Static RDP (Required Deceleration Parameter) calculates the required deceleration for the vehicle to stop safely given its current speed and position on the road. RDP is computed as $RDP = \frac{v^2}{2 \times r \times g}$, where r is distance to the crossing line at the intersection, and v is the current speed of the vehicle or other ground transportation entity. g is the gravity acceleration constant.

- 5 A vehicle is classified as dangerous (that is, the vehicle has or will create a dangerous situation) if its required deceleration is larger than the chosen RDP threshold RDP_{warn} . In practice, a

vehicle is classified as dangerous if at any time, $r < r_{alert}$, where $r_{alert} = \frac{v^2}{2 \times RDP_{alert}}$.

Similar to the static TTI algorithm, the RDP_{alert} parameter reflects the conservativeness level of the rule-based algorithm.

- 10 We use rule-based approaches as a baseline for the assessment of the performance of our machine learning algorithms, and in some instances, we run them in parallel to the machine learning algorithms to capture the rare cases in which machine learning might not able to predict.

Machine Learning

- Modeling driver's behaviors have been shown to be a complex task given the complexity of human behavior. See H. M. Mandalia and D. D. Dalvucci, "Using Support Vector Machines for Lane-Change Detection," Human Factors and Ergonomics Society Annual Meeting Proceedings, vol. 49, pp. 1965-1969, 2005. Machine learning techniques are well suited to model human behavior but need to "learn" using training data to work properly. To provide superior detection and prediction results, we use machine learning to model traffic detected at an intersection or other features of a ground transportation network during a training period before the alerting process is applied to current traffic during a deployment phase. Machine learning can be used also to model driver responses using in-vehicle data from onboard equipment (OBE), and could also be based on in-vehicle sensors and history of driving records and preferences. We also use machine learning models to detect and predict vulnerable road user (e.g., pedestrian) trajectories, behaviors and intents. Machine learning can be used also to model vulnerable road users responses from on-person equipment (OPE). These models could include interactions between entities, vulnerable road users, and between one or multiple entities and one or multiple vulnerable road users.
- 15
20
25

Machine learning techniques could be also used to model the behaviors of non-autonomous ground transport entities. By observing or communicating or both with a non-autonomous ground transportation entity, machine learning can be used to predict its intent and communicate with it and with other involved entities when a near-miss or accident or other dangerous situation is predicted.

The machine learning mechanism works in of two phases: 1) training and 2) deployment.

Training phase

After installation, the RSE starts collecting data from the sensors to which it has access. Since AI model training requires intense computational capacity, it is usually performed on powerful servers that have multiple parallel processing modules to speed up the training phase. For this reason, the data acquired at the location of the RSE on the ground transportation network can be packaged and sent to a remote powerful server shortly after the acquisition. This is done using an Internet connection. The data is then prepared either automatically or with the help of a data scientist. The AI model is then built to capture important characteristics of the flow of traffic of vehicles and other ground transportation entities for that intersection or other aspects of the ground transportation network. Captured data features may include location, direction, and movement of the vehicles or other ground transportation entities, which can then be translated to intent and behavior. Knowing intent, we can predict actions and future behavior of vehicles or other ground transportation entities approaching the traffic location using the AI model, with high accuracy. The trained AI model is tested on a subset of the data that has not been included in the training phase. If the performance of the AI model meets expectations, the training is considered complete. This phase is repeated iteratively using different model parameters until a satisfactory performance of the model is achieved.

Deployment phase

In some implementations, the complete tested AI model is then transferred through the Internet to the RSE at the traffic location in the ground transportation network. The RSE is then ready to process new sensor data and perform prediction and detection of dangerous situations such as traffic light violations. When a dangerous situation is predicted, the RSE will generate an appropriate alert message. The dangerous situation can be predicted, the alert message generated,

and the alert message broadcast to and received by vehicles and other ground transportation entities in the vicinity of the RSE before the predicted dangerous situation occurs. This allows the operators of the vehicles or other ground transportation entities ample time to react and engage in collision avoidance. The outputs of the AI models from the various intersections at which the corresponding RSEs are located can be recorded and made available online in a dashboard that incorporates all the data generated and displayed in an intuitive and user-friendly manner. Such a dashboard could be used as an interface with the customer of the system (e.g., a city traffic engineer or planner). One example of dashboard is a map with markers that indicate the locations of the monitored intersections, violation events that have occurred, statistics and analytics based on the AI predictions and actual outcomes.

Smart RSE (SRSE) and the connected entity/non-connected entity bridge

As suggested earlier, there is a gap between the capabilities and actions of connected entities and non-connected entities. For example, connected entities are typically cooperative entities that continuously advertise to the world their location and safety system status such as speed, heading, brake status, and headlight status. Non-connected entities are not able to cooperate and communicate in these ways. Therefore, even a connected entity will be unaware of a non-connected entity that is not in the connected entity's vicinity or out of sensor range due to interference, distance, or the lack of a good vantage point.

With the proper equipment and configuration, RSEs can be made capable of detecting all entities using the ground transportation network in their vicinities, including non-connected entities. Specialized sensors may be used to detect different types of entities. For example, radars are suitable for detecting moving metallic objects such as cars, buses and trucks. Such road entities are most likely moving in a single direction towards the intersection. Cameras are suitable of detecting vulnerable road users who may wander around the intersection looking for a safe time to cross.

Placing sensors on components of the ground transportation network has at least the following advantages:

- Good vantage point: Infrastructure poles, beams, and support cables usually have an elevated vantage point. The elevated vantage points allow for a more general view of the intersection.

This is like an observation tower at an airport where controllers have a full view of most of the important and vulnerable users on the ground. For ground transportation entities, by contrast, the views from the vantage point of sensors (camera, lidar, radar, etc....or others) can be obstructed or disrupted by a truck in a neighboring lane, direct sunlight, or other interference. The sensors at
5 the intersection can be chosen to be immune or less susceptible to such interference. A radar, for example, is not affected by sunlight and will remain effective during the evening commute. A thermal camera will be more likely to detect a pedestrian in a bright light situation where the view of an optical camera becomes hindered.

- Fixed location: Sensors situated at the intersection can be adjusted and fixed to sense in a
10 specific direction that can be optimal for detecting important targets. This will help the processing software to better detect objects. As an example, if a camera has a fixed view, the background (non-moving objects and structures) information in the fixed view can be easily detected and used to improve the identification and classification of relatively important moving entities.
- 15 Fixed sensor location also enables easier placement of every entity in a unified global view of the intersection. Since the sensor view is fixed, the measurements from the sensor can be easily mapped to a unified global location map of the intersection. Such a unified map is useful when performing global analysis of traffic movements from all directions to study the interactions and dependencies of one traffic flow on another. An example would be in detecting a near miss
20 (dangerous situation) before it happens. When two entities are traveling along intersecting paths, a global and unified view of the intersection will enable the calculation of the time of arrival of each entity to the point of intersection of the respective paths. If the time is within a certain limit or tolerance, a near miss may be flagged (e.g., made the subject of an alert message) before it happens.
- 25 With the help of the sensors that are installed on components of the infrastructure, smart RSEs (SRSEs) can bridge this gap and allow connected entities to be aware of “dark” or non-connected entities.

Figure 8 depicts a scenario that explains how strategically placed sensors can help connected entities identify the speed and location of non-connected entities.

A connected entity 1001, is traveling along a path 1007. The entity 1001 has a green light 1010. A non-connected entity 1002 is traveling along a path 1006. It has a red light 1009 but will be making a right on red along path 1006. This will place it directly in the path of the entity 1001. A dangerous situation is imminent since the entity 1001 is unaware of the entity 1002. Because the entity 1002 is a non-connected entity it is unable to broadcast (e.g., advertise) its position and heading to other entities sharing the intersection. Moreover, the entity 1001, even though it is connected, is unable to "see" the entity 1002 which is obscured by the building 1008. There is a risk of the entity 1001 going straight through the intersection and hitting the entity 1002.

If the intersection is configured as a smart intersection, a radar 1004 mounted on a beam 1005 above the road at the intersection will detect the entity 1002 and its speed and distance. This information can be relayed to the connected entity 1001 through the SRSE 1011 serving as a bridge between the non-connected entity 1002 and the connected entity 1001.

Artificial Intelligence and Machine Learning

Smart RSEs also rely on learning traffic patterns and entity behaviors to better predict and prevent dangerous situations and avoid collisions. As shown in figure 8, the radar 1004 is always sensing and providing data for every entity moving along approach 1012. This data is collected and transferred to the cloud, either directly or through an RSE, for example, for analysis and for building and training a model that closely represents the traffic along approach 1012. When the model is complete, it is downloaded to the SRSE 1011. This model can then be applied to every entity moving along approach 1012. If an entity is classified by the model as one that is (or is going to) violate the traffic rules, a warning (alert) may be broadcast by the SRSE to all connected entities in the vicinity. This warning, known as intersection collision avoidance warning, will be received by the connected entities and can be acted upon to take account of the dangerous situation and avoid a collision. With the proper traffic model, a violating entity can be detected in advance, giving connected entities using the intersection enough time to react and avoid a dangerous situation.

With the help of multiple sensors (some mounted high on components of the infrastructure of the ground transportation network), artificial intelligence models, and accurate traffic models, an SRSE can have a virtual overview of the ground transportation network and be aware of every

entity within its field of view including non-connected entities in the field of view that are not “visible” to connected entities in the field of view. The SRSE can use this data to feed the AI model and provide alerts to connected entities on behalf of non-connected entities. A connected entity would not otherwise know that there are non-connected entities sharing the road.

- 5 SRSEs have high power computing available at the location of the SRSE either within the same housing or by connection to a nearby unit or through the Internet to servers. An SRSE can process data received directly from sensors, or data received in broadcasts from nearby SRSEs, emergency and weather information, and other data. An SRSE is also equipped with high capacity storage to aid in storing and processing data. High bandwidth connectivity is also
- 10 needed to help in transferring raw data and AI models between the SRSE and even more powerful remote servers. SRSEs enhance other traffic hazard detection techniques using AI to achieve high accuracy and provide additional time to react and avoid a collision.

SRSEs can remain compatible with current and new standardized communication protocols and, therefore, they can be seamlessly interfaced with equipment already deployed in the field.

- 15 SRSEs can also reduce network congestion by sending messages only when necessary.

Global and unified intersection topology

Effective traffic monitoring and control of an intersection benefits from a bird’s eye view of the intersection that is not hindered by obstacles, lighting, or any other interference.

As discussed above, different types of sensors can be used to detect different types of entities.

- 20 The information from these sensors can be different, e.g., inconsistent with respect to the location or motion parameters that its data represents or the native format of the data or both. For example, radar data typically includes speed, distance, and maybe additional information such as the number of moving and stationary entities that are in the field of view of the radar. Camera data, by contrast, can represent an image of the field of view at any moment in time. Lidar data
- 25 may provide the locations of points in 3D space that correspond to the points of reflection of the laser beam emitted from the lidar at a specific time and heading. In general, each sensor provides data in a native format that closely represents the physical quantities they measure.

To get a unified view (representation) of the intersection, fusion of data from different types of sensors is useful. For purposes of fusion, the data from various sensor is translated into a

common (unified) format that is independent of the sensor used. The data included in the unified format from all of the sensors will include the global location, speed, and heading of every entity using the intersection independently of how it was detected.

Armed with this unified global data, a smart RSE can not only detect and predict the movement of entities, but also can determine the relative positions and headings of different entities with respect to each other. Therefore, the SRSE can achieve improved detection and prediction of dangerous situations.

For example, in the scenario shown in figure 9, a motorized entity 2001 and a vulnerable road user 2002 share the same pedestrian crossing. The entity 2001 is traveling along a road 2007 and is detected by radar 2003. The vulnerable road user 2002 walking along sidewalk 2006 is detected by a camera 2004. The vulnerable road user 2002 may decide to cross the road 2007 using a crosswalk 2005. Doing so places the road user 2002 in the path of entity 2001 creating a possible dangerous situation. If the data from each of the sensors 2003 and 2004 were considered independently and no other information were considered, the dangerous situation would not be identified since each of the sensors can only detect the entities in its respective fields of view. Additionally, each of the sensors may not be able to detect objects that they are not designed to detect. However, when a unified view is considered by the SRSE, the locations and dynamics of the entity 2001 and of the vulnerable road user 2002 can be placed in the same reference frame: a geographic coordinate system such as a map projections or other coordinate system. When considered within a common reference system, the fused data from the sensors can be used to detect and predict a dangerous situation that may arise between the two entities 2001 and 2002. We will discuss the translation between the sensor space and the unified space in the following paragraphs.

Radar data to unified reference translation

As shown in figure 10, a radar 3001 is used to monitor road entities traveling along a road having two lanes 3005 and 3008 with centerlines 3006 and 3007 respectively. A stop bar 3003 indicates the end of lanes 3005 and 3008. T 3006 can be defined by a set of markers 3003 and 3004. Figure 10 shows only two markers but, in general, the centerline is a piecewise linear function. The global locations of markers 3003 and 3004 (and the other markers, not shown) are

predefined by the design of the roadway and are known to the system. The precise global location of radar 3001 can also be determined. Distances 3009 and 3010 of markers 3003 and 3004 from the radar 3001 can, therefore, be calculated. The distance 3011 of the entity 3002 from the radar 3001 can be measured by the radar 3001. Using simple geometry, the system can
5 determine the location of the entity 3002 using the measured distance 3011. The result is a global location since it is derived from the global locations of markers 3003, 3004 and the radar 3001. Since every roadway can be approximated by a generalized piecewise linear function, the method above can be applied to any roadway that can be monitored by a radar.

Figure 11 shows a similar scenario on a curved road. Radar 4001 monitors entities moving along
10 a road 4008. The markers 4003 and 4004 represent a linear segment 4009 (of the piecewise linear function) of the centerline 4007. The distances 4005 and 4006 represent the normal distance between the plane 4010 of the radar 4001 and the markers 4003 and 4004 respectively. Distance 4007 is the measured distance of entity 4002 from the radar plane 4010. Following the discussion above, given the global locations of the radar 4001 and the markers 4003 and 4004,
15 the global location of the entity 4002 can be calculated using simple ratio arithmetic.

Camera data to unified reference translation

Knowing the height, global location, direction, tilt and field of view of a camera, calculating the global location of every pixel in the camera image become straight forward using existing 3D geometry rules and transformations. Consequently, when an object is identified in the image, its
20 global location can be readily deduced by knowing the pixels it occupies. It is beneficial to note that the type of camera is irrelevant if its specifications are known, such as sensor size, focal length, or field of view, or combinations of them.

Figure 12 shows a side view of a camera 5001 looking at an entity 5002. The height 5008 and tilt angle 5006 of the camera 5001 can be determined at the time of installation. The field of view
25 5007 can be obtained from the specifications of the camera 5001. The global location of the camera 5001 can also be determined at the time of installation. From the known information, the system can determine the global positions of the points 5003 and 5004. The distance between points 5003 and 5004 is also divided into pixels on the image created by the camera 5001. This number of pixels is known from the camera 5001 specifications. The pixels occupied by the

entity 5002 can be determined. The distance 5005 can therefore be calculated. The global location of entity 5002 can also be calculated.

A global unified view of any intersection can be pieced together by fusing the information from various sensors. Figure 13 depicts a top view of a four-way intersection. Every leg of the intersection is divided by a median 6003. The intersection in the figure is being monitored by two different types of sensors, radars and cameras, and the principles discussed here can be generalized to other types of sensors. In this example, radar monitored regions 6001 overlap camera monitored regions 6002. With a unified global view, every entity that travels between regions will remain tracked within the unified global view. This makes determinations by the SRSE, for example, of relationships between the motions of different entities easily possible. Such information will allow for a truly universal birds eye view of the intersection and roadways. The unified data from the sensors can then be fed into artificial intelligence programs as described in the following paragraphs.

Figure 2, discussed above illustrated components of an RSE. In addition, in an SRSE, the processing unit may also include one or several specialized processing units that can process data in parallel. An example of such units are graphic processing units or GPUs. With the aid of GPUs or similar hardware, machine learning algorithms can run much more efficiently at the SRSE and will be able to provide results in real time. Such a processing architecture enables real time prediction of dangerous situations and therefore enables sending warnings early on to allow the entities enough time to react and avoid collisions. In addition, because an SRSE can run processes that can use the data from different sensors and different types of sensors, the SRSE can build a unified view of the intersection that would help in the analysis of traffic flows and the detection and prediction of dangerous situations.

Use Cases

A wide variety of cases can benefit from the system and the early warnings that it can provide for collision avoidance. Examples are provided here.

Case 1: Vulnerable ground transportation entities

As shown in figure 4, a roadway that crosses a typical intersection 409 may have a pedestrian crosswalk including specific crossing areas 401, 402, 403, 404 that pedestrians and other

vulnerable road users (vulnerable road users) may use to walk across the roadway. Sensors that are adequate to detect such crossings or other vulnerable users are placed at one or more vantage points that allow the monitoring of the crosswalk and its surroundings. During a training phase, the collected data can be used to train an artificial intelligence model to learn about the behavior of vulnerable road users at the intersection. During a deployment phase, the AI model then can use current data about a vulnerable road user to predict, for example, that the vulnerable road user is about to cross the roadway, and to make that prediction before the vulnerable road user begins to cross. When behavior and intent of pedestrians and other vulnerable road users, drivers, vehicles, and other people and ground transportation entities can be predicted in advance, early warnings (e.g., alerts) can be sent to any or all of them. Early warning can enable vehicles to stop, slow down, change paths, or combinations of them, and can enable vulnerable road users to refrain from crossing the road when a dangerous situation is predicted to be imminent.

In general, sensors are used to monitor all areas of possible movement of vulnerable road users and vehicles in the vicinity of an intersection. The types of sensors used depend on the types of subjects being monitored and tracked. Some sensors are better at tracking people and bicycles or other non-motorized vehicles. Some sensors are better at monitoring and tracking motorized vehicles. The solution described here is sensor and hardware agnostic, because the type of sensor is irrelevant if it provides appropriate data at a sufficient data rate which can be depend on the types of subjects being monitored and tracked. For example, Doppler radar would be an appropriate sensor to monitor and track the speed and distance of vehicles. The data rate, or sampling rate, is the rate at which the radar is able to provide successive new data values. The data rate must be fast enough to capture the dynamics of the motions of the subject being monitored and tracked. The higher the sampling rate, the more details are captured and the more robust and accurate the representation of the motion by the data becomes. If the sampling rate is too low, and the vehicle travels a significant distance between two sample instances, it becomes difficult to model the behavior because of the missed details during the intervals for which data is not generated.

For a pedestrian crossing, sensors will monitor the pedestrian and other vulnerable road users (e.g., cyclists) crossing at the intersection and the areas in the vicinity of the intersection. The data from these sensors may be segmented as representing conditions with respective different

virtual zones to help in detection and localization. The zones can be chosen to correspond to respective critical areas where dangerous situations may be expected, such as sidewalks, entrances of walkways, and incoming approaches 405, 406, 407, 408 of the roads to the intersection. The activity and other conditions in every zone is recorded. Records can include, but are not limited to kinematics (e.g., location, heading, speed, and, acceleration) and facial and body features (e.g., eyes, posture)

The number of sensors, number of zones, and shapes of zones are specific to every intersection and to every approach to the intersection.

Figure 5 depicts a plan view of a typical example setup showing different zones used to monitor and track the movement and behavior of pedestrians or other vulnerable road users, and motorized and non-motorized vehicles and other ground transportation entities.

Sensors are set up to monitor a pedestrian crosswalk across a road. Virtual zones (301, 302) may be placed on the sidewalks and along the crosswalk. Other sensors are placed to monitor vehicles and other ground transportation entities proceeding on the road leading to the crosswalk, and virtual zones (303, 304) are strategically placed to aid in detecting incoming vehicles and other ground transportation entities, their distances from the crosswalk, and their speeds, for example.

The system (e.g., the RSE or SRSE associated with the sensors) collects streams of data from all sensors. When the system is first put into operation, to help with equipment calibration and functionality, an initial rule-based model may be deployed. In the meantime, sensor data (e.g., speed and distance from radar units, images and video from cameras) is collected and stored locally at the RSE in preparation, in some implementations, to be transferred to a remote computer that is powerful enough to build an AI model of the behavior of the different entities of the intersection using this collected data. In some cases, the RSE is a SRSE capable of generating the AI model itself.

The data is then prepared, and trajectories are built for every ground transportation entity passing through the intersection. For example, trajectories can be extracted from radar data by stitching together points of different distances that belong to the same entity. Pedestrian trajectories and behavior can be, for example, extracted from camera and video recordings. By performing video

and image processing techniques, the movement of the pedestrian can be detected in images and videos and their respective trajectories can be deduced.

For human behavior, an intelligent machine learning based model typically outperforms a simple rule based on simple physics. This is because human intent is difficult to capture, and large
5 datasets are needed to be able to detect patterns.

When the machine learning (AI) model is completed at the server, it is downloaded to the RSE through the Internet, for example. The RSE then applies current data captured from the sensors to the AI model to cause it to predict intent and behavior, to determine when a dangerous situation is imminent, and to trigger corresponding alerts that are distributed (e.g., broadcast) to
10 the vehicles and other ground transportation entities and to the vulnerable road users and drivers as early warnings in time to enable the vulnerable road users and drivers to undertake collision avoidance steps.

This example setup can be combined with any other use case, such as traffic at signalized intersections or level crossings.

15 Case 2: Signalized intersection

In the case of a signalized intersection (e.g., one controlled by a traffic light) the overall setup of the system is done as in case 1. One difference may be the types of sensors used to monitor or track vehicle speed, heading, distance, and location. The setup for the pedestrian crossing of case
1 can also be combined with the signalized intersection setup for a more general solution.

20 The concept of operations for the signalized intersection use case is to track road users around the intersection using external sensors collecting data about the users or data communicated by the users themselves, predict their behaviors and broadcast alerts through different communication means about upcoming hazardous situations, generally due to violations of intersection traffic rules, such as violating a red-light signal.

25 Data on road users can be collected using (a) entity data broadcast by each entity itself about its current state, through a BSM or a PSM for instance; and (b) sensors installed externally on infrastructure or on vehicles, such as doppler radars, ultrasonic sensors, vision or thermal cameras, lidars, and others. As mentioned earlier, the type of sensor selected and its position and orientation at the intersection should provide the most comprehensive coverage of the

intersection, or the part of it under study and that the data collected about the entities approaching the intersection is the most accurate. Thus, the data collected will allow reconstruction of the current states of road users and creation of an accurate, timely, useful VBSM (virtual basic safety message) or VPSM (virtual personal safety message). The frequency at which data should be collected depends on the potential hazard of each type of road user and the criticality of a potential violation. For instance, motorized vehicles traveling at high speeds in the intersection usually require data updates 10 times per second to achieve real time collision avoidance; pedestrians crossing the intersection at much lower speeds can require data updates as low as 1 time per second.

- 10 As noted earlier, figure 4 depicts an example of a signalized intersection plan view with detection virtual zones. These zones can segment every approach to the intersection into separate lanes 410, 411, 412, 413, 405, 406, 407, 408 and may also separate each lane into areas that correspond to general ranges of distance from the stop bar. The choice of these zones may be performed empirically to match the character of the specific approaches and intersection in general. Segmenting the intersection allows for more accurate determinations of relative heading, speed, acceleration, and positioning for each road user and in turn a better assessment of the potential hazard that road user presents to other ground transportation entities.

- In order to determine whether an observed traffic situation is a dangerous situation, the system also needs to compare the outcome of the predicted situation with the traffic light state and account for local traffic rules (e.g. left-turn lanes, right-turn on red, and others). Therefore, it is necessary to collect and use the intersection's signal phase and timing (SPaT) information. SPaT data can be collected by interfacing directly with the traffic light controller at the intersection, generally through a wired connection reading the data, or by interfacing with the traffic management system to receive the required data, for instance through an API. It is important to collect SPaT data at a rate as close as possible to the rate at which road user data is collected to ensure that road user state is always synchronized with traffic signal state. An added complexity to the requirement of knowing SPaT information is that modern traffic control strategies employed to regulate traffic flow around intersections are not based on fixed timings and use algorithms that can dynamically adapt to real-time traffic conditions. It is thus important to incorporate SPaT data prediction algorithms to insure the highest accuracy in violation

prediction. These SPaT data prediction algorithms can be developed using rule-based methods or machine learning methods.

For each approach to the intersection, data is collected by the RSE (or SRSE) and a machine learning (AI) model is constructed to describe the behavior of the vehicles corresponding to the collected data. Current data collected at the intersection is then applied to the AI model to produce an early prediction whether a vehicle or other ground transportation entity traveling on one of the approaches to the intersection is, for example, about to violate the traffic light. If a violation is imminent, a message is relayed (e.g., broadcast) from the RSE to ground transportation entities in the vicinity. Vehicles (including the violating vehicle) and pedestrians or other vulnerable road users will receive the message and have time to take appropriate pre-emptive measures to avoid a collision. The message can be delivered to the ground transportation entities in one or more of the following ways, among others: a blinking light, sign, or radio signal.

If a vehicle or other entity approaching the intersection is equipped with an OBE or an OPE, it will be able to receive the message broadcast from the RSE that a potential hazard has been predicted at the intersection. This allows the user to be warned and to take appropriate pre-emptive measures to avoid a collision. If the violating road user at the intersection is also equipped with an OBE or an OPE, the user will also receive the broadcast alert. Algorithms on the OBE or an OPE can then reconcile the message with the violating behavior of the user and warn the user adequately.

The decision to send an alert is dependent not only on the vehicle behavior represented by the data collected by the sensors at the intersection. Although the sensors play a major role in the decision, other inputs are also considered. These inputs may include, but not be limited to, information from a nearby intersection (if a vehicle ran the light at a nearby intersection, there is higher probability that it would do the same at this intersection), information from other cooperative vehicles, or even the vehicle itself, if for example it is reporting that it has a malfunction.

Case 3: Non-signalized intersection

Non-signalized controlled intersections, such as a stop sign or yield sign-controlled intersection, can be monitored as well. Sensors are used to monitor the approach controlled by the traffic sign and predictions can be made about incoming vehicles, similar to predictions about incoming vehicles on an approach to a signalized intersection. The rules of the roads at non-signalized controlled intersections are typically well defined. The ground transportation entity on an approach controlled by a stop sign must come to a full stop. In a multi-way stop intersection, the right-of-way is determined by the order the ground transportation entities reach the intersection. A special case can be considered with a one-way stop. A set of sensors can monitor the approach that does not have a stop sign as well. Such a setup can assist in stop sign gap negotiations. For a yield sign controlled intersection, a ground transportation entity on an approach controlled by a yield sign must reduce its speed to give right-of-way to other ground transportation entities in the intersection.

A main challenge is that due to internal (e.g., driver distraction) or external (e.g., lack of visibility) factors, ground transportation entities violate the rules of the road, and put other ground transportation entities at risk.

In the general case of stop-sign controlled intersections (i.e., each approach is controlled by a stop sign), the overall setup of the system is done as in case 1. One difference may be the types of sensors used to monitor or track vehicle speed, heading, distance, and location. Another difference is the lack of traffic light controllers with the rules of the roads being indicated by the road signs. The setup for the pedestrian crossing of case 1 can also be combined with the non-signalized controlled intersection setup for a more general solution.

Figure 4 could also be understood to depict an example of a four-way stop intersection plan view with detection virtual zones. These zones can segment every approach to the intersection into separate lanes 410, 411, 412, 413, 405, 406, 407, 408 and may also separate each lane into areas that correspond to general ranges of distance from the stop bar. The choice of these zones may be made empirically to match the character of the specific approaches and the intersection in general.

In a manner similar to the one described above for figure 4, current data collected at the intersection is applied to the AI model to produce an early prediction whether a vehicle or other

ground transportation entity traveling on one of the approaches to the intersection is about to violate the stop sign. If a violation is imminent, messages can be handled similarly to the previously described case involving a traffic light violation.

Also similarly to the previous description, the decision to send an alert can be based on factors
5 described previously and on other information such as whether the vehicle ran the stop sign at a nearby intersection, suggesting a higher probability that it would do the same at this intersection).

Figure 18 illustrates a use case for controlled non-signalized intersection. It explains how the SRSE with strategically placed sensors can warn a connected entity of an impending dangerous
10 situation arising from a non-connected entity.

A connected entity 9106, is traveling along a path 9109. The entity 9106 has the right of way. A non-connected entity 9107 is traveling along path 9110. The entity 9107 has a yield sign 9104 and will be merging onto the path 9109 without giving right of way to the entity 9106 placing it directly in the path of the entity 9106. A dangerous situation is imminent since the entity 9106 is
15 unaware of the entity 9107. Because the entity 9107 is a non-connected entity, it is unable to advertise (broadcast) its position and heading to other entities sharing the intersection. Moreover, the entity 9106 may not be able to “see” the entity 9107 which is not in its direct field of view. If the entity 9106 proceeds along its path it may eventually have a collision with the entity 9107.

Because the intersection is a smart intersection, a radar 9111 mounted on a beam 9102 above the
20 road will detect the entity 9107. It will also detect the entity 9107 speed and distance. This information can be relayed as an alert to the connected entity 9106 through the SRSE 9101. The SRSE 9101 has a machine learning model for entities moving along the approach 9110. The entity 9107 will be classified by the model as a potential violator of the traffic rule, and a warning (alert) will be broadcast to the connected entity 9106. This warning is sent in advance
25 giving the entity 9106 enough time to react and prevent a dangerous situation.

Case 4: Level crossings

Level crossings are dangerous because they may carry motorized vehicles, pedestrians, and rail vehicles. In many cases, the road leading to the level crossing falls in the blind spot of an operator (e.g., conductor) of a train or other rail vehicle. Since rail vehicle drivers operate mainly

on line-of-sight information, this increases the possibility of an accident if the road user violates the rail vehicle's right of way and crosses a level crossing when it is not permitted to cross.

The operation of the level crossings use case is similar to the signalized intersection use case, in the sense that a level crossing is a conflict point between road and rail traffic often regulated by traffic rules and signals. Therefore, this use case also requires collision avoidance warnings to increase safety around level crossings. Rail traffic can have a systematic segregated right of way, e.g., high-speed rail, or no segregated right of way, e.g., light urban rail or streetcars. With light rail and streetcars, the use case becomes even more important since these rail vehicles also operate on live roads and have to follow the same traffic rules as road users.

Figure 6 depicts a general use of a level crossing where a road and a pedestrian crossing cross a railroad. Similar to the pedestrian crossing use case, sensors are placed to collect data on the movement and intent of pedestrians. Other sensors are used to monitor and predict movement of road vehicles that are set to approach the crossing. Data on road users can also be collected from road user broadcasts (e.g., BSMs or PSMs). Data from nearby intersections, vehicles, and remote command and control centers may be used in the decision to trigger an alert.

Data on SPaT for road and rail approaches will also need to be collected in order to adequately assess the potential for a violation.

Similarly to the signalized intersection use case, the data collected enables the creation of predictive models using rule-based and machine learning algorithms.

In this use case, the rail vehicle is equipped with an OBE or an OPE in order to receive collision avoidance warnings. When a violation of the rail vehicle's right of way is predicted, the RSE will broadcast an alert message, warning the rail vehicle driver that a road user is in its intended path and allowing the rail vehicle driver to take pre-emptive actions with enough time to avoid the collision.

If the violating road user is also equipped with an OBE or an OPE, the message broadcast by the RSE will also be received by the violating road user. Algorithms on the OBE or an OPE can then reconcile the received message with the violating behavior of the user and warn the user adequately.

Virtual connected ground transportation environment (bridging the gap)

As discussed above, a useful application of the system is to create a virtual connected environment on behalf of non-connected ground transportation entities. An impediment to the adoption of connected technology is not only the absence of infrastructure installations, but also the almost non-existence of connected vehicles, connected vulnerable road users, and connected
5 other ground transportation entities.

With respect to connected vehicles, in some regulatory regimes, such vehicles are always sending what are called basic safety messages (BSMs). BSMs contain, among other information, the location, heading, speed, and future path of the vehicle. Other connected vehicles can tune in to these messages and use them to create a map of vehicles present in their surroundings.
10 Knowing where the surrounding vehicles are, a vehicle, whether it is autonomous or not, will have information useful to maintain a high level of safety. For example, an autonomous vehicle can avoid making a maneuver if there is a connected vehicle in its path. Similarly, a driver can be alerted if there is some other vehicle in the path that he is planning to follow such as a sudden lane change.

15 Until all ground transportation entities are equipped to send and receive traffic safety messages and information, some road entities will be “dark” or invisible to the rest of the road entities. Dark road entities pose a risk of a dangerous situation.

Dark road entities do not advertise (e.g., broadcast) their location, so they are invisible to connected entities that may expect all road entities to broadcast their information (that is, to be
20 connected entities). Although onboard sensors can detect obstacles and other road entities, the ranges of these sensors tend to be too short to be effective in preventing dangerous situations and collisions. Therefore, there is a gap between the connectivity of connected vehicles and the lack of connectivity of non-connected vehicles. The technology described below is aimed to bridge this gap by using intelligence on the infrastructure that can detect all vehicles at the intersection
25 or other component of the ground transportation network and send messages on behalf of non-connected vehicles.

The system can establish a virtual connected ground transportation environment, for example, at an intersection, that can bridge the gap between the future when most vehicles (and other ground transportation entities) are expected to be connected entities and the current time when most

vehicles and other ground transportation entities have no connectivity. In the virtual connected ground transportation environment, smart traffic lights and other infrastructure installations can use sensors to track all vehicles and other ground transportation entities (connected, non-connected, semi-autonomous, autonomous, non-autonomous) and (in the case of vehicles)
5 generate virtual BSM messages (VBSM) on their behalf.

A VBSM message can be considered a subset of a BSM. It may not contain all the fields required to create a BSM but can contain all the localization information including location, heading, speed and trajectory. Since V2X communication is standardized and anonymized, VBSM and BSM cannot be differentiated easily and follow the same message structure. The
10 main difference between the two messages is the availability of the sources of the information populating these messages. A VBSM might lack data and information not easily generated by external sensors such as steering wheel angle, brake status, tire pressure or wiper activation.

With the proper sensors installed, an intersection with smart RSE can detect all the road entities that are travelling through the intersection. The SRSE can also transform all data from multiple
15 sensors into a global unified coordinate system. This global unified system is represented by the geographical location, speed and heading of every road entity. Every road entity, whether it is connected or not, is detected by the intersection equipment and a global unified location is generated on its behalf. Standard safety messages can, therefore, be broadcast on behalf of the road entities. However, if the RSE broadcasts a safety message for all entities it detects, it may
20 send a message on behalf of a connected road entity. To address the conflict, the RSE can filter the connected road entities from its list of dark entities. This can be achieved because the RSE is continuously receiving safety messages from connected vehicles, and the RSE sensors are continuously detecting road entities passing through the intersection. If the location of a detected road entity matches a location that from which a safety message is received by the RSE receiver,
25 the road entity is assumed to be a connected and no safety message is broadcast on its behalf by the RSE. This is depicted in figure 15.

By creating the bridge between connected and non-connected vehicles, connected entities (including autonomous vehicles) can safely maneuver through intersections with complete awareness of all the road entities nearby.

This aspect of the technology is illustrated in Figure 17. An intersection 9001 has multiple road entities at a given time. Some of these entities are non-connected 9004, 9006 and others are connected 9005, 9007. Vulnerable road users 9004, 9007 are detected by a camera 9002.

Motorized road entities 9005, 9006 are detected by radars 9003. The location of each road entity is calculated. Broadcasts from connected road entities are also received by the RSE 9008. The locations of entities from which messages are received are compared with the locations at which entities are detected. If two entities match within a predetermined tolerance, the entity at that location is considered connected and no safety message is sent on its behalf. The rest of the road entities that have no matching received location are considered dark. Safety messages are broadcast on their behalf.

For collision warnings and intersection violation warnings that are an integral part of V2X protocols, every entity needs to be connected for the system to be effective. That requirement is a hurdle in the deployment of V2X devices and systems. Intersections equipped with smart RSE will address that concern by providing a virtual bridge between connected and non-connected vehicles.

The US DOT (Department of Transportation) and NHTSA (National Highway Traffic Safety Administration) identify a number of connected vehicle applications that will use BSMS and help substantially decrease non-impaired crashes and fatalities. These applications include, but are not limited to, Forward Collision Warning (FCW), Intersection Movement Assist (IMA), Left Turn Assist (LTA), Do Not Pass Warning (DNPW), and Blind Spot/Lane Change Warning (BS/LCW). The US DOT and NHTSA define these applications as follows.

An FCW addresses rear-end crashes and warns drivers of stopped, slowing, or slower vehicles ahead. An IMA is designed to avoid intersection crossing crashes and warns drivers of vehicles approaching from a lateral direction at an intersection covering two major scenarios: Turn-into path into same direction or opposite direction and straight crossing paths. An LTA addresses crashes where one involved vehicle was making a left turn at the intersection and the other vehicle was traveling straight from the opposite direction and warns drivers to the presence of oncoming, opposite-direction traffic when attempting a left turn. A DNPW assists drivers to avoid opposite-direction crashes that result from passing maneuvers and warns a driver of an oncoming, opposite-direction vehicle when attempting to pass a slower vehicle on an undivided

two-lane roadway. A BS/LCW addresses crashes where a vehicle made a lane changing/merging maneuver prior to the crashes and alerts drivers to the presence of vehicles approaching or in their blind spot in the adjacent lane.

V2X protocols stipulate that these applications should be achieved using vehicle-to-vehicle (V2V) communications, where one connected remote vehicle would broadcast basic safety messages to a connected host vehicle. The host vehicle's OBE would in turn try to reconcile these BSMs with its own vehicle parameters, such as speed, heading and trajectory and determine if there is a potential danger or threat presented by the remote vehicle as described in the applications above. Also, an autonomous vehicle will benefit specifically from such an application, since it allows surrounding vehicles to communicate intent, which is a key piece of information not contained in the data collected from its onboard sensors.

However, today's vehicles are not connected and, as mentioned earlier, it will take a significant period until the proportion of connected vehicles is high for BSMs to work as explained above. Therefore, in an environment in which the proportion of connected vehicles is small, the connected vehicles are not required to receive and analyze the large number of BSMs they would otherwise receive in an environment having a proportion of connected vehicles large enough to enable the applications described above and benefit fully from V2X communication.

VBSMs can help bridge the gap between the current environment having largely unconnected entities and a future environment having largely connected entities and enable the applications described above, during the interim.. In the technology that we describe here, a connected vehicle receiving a VBSM will process it as a regular BSM in the applications. Since VBSMs and BSMs follow the same message structure and VBS Ms contain substantially the same basic information as a BSM, e.g., speed, acceleration, heading, past and predicted trajectory, the outcome of applying the messages to a given application will be substantially the same.

For example, consider an intersection with non-protected left turns, where the connected host vehicle is about to attempt a left turn at a moment when an unconnected remote vehicle is traveling straight from the opposing direction with right of way. This is a situation where completion of the maneuver depends on the host vehicle driver's judgment of the situation. A wrong assessment of the situation may result in a conflict and a potential near-collision or

collision. External sensors installed on the surrounding infrastructure can detect and track the remote vehicle or even both vehicles, collect basic information such as speed, acceleration, heading and past trajectory and transmit them to the RSE, which can in turn build the predicted trajectory for the remote vehicle using rule-based or machine learning algorithms or both, 5 populate the required fields for the VBSM and broadcast it on behalf of the unconnected remote vehicle. The host vehicle's OBE will receive the VBSM with information about the remote vehicle and process it in its LTA application to determine whether the driver's maneuver presents a potential danger and if the OBE should display a warning to the host vehicle's driver to take preemptive or corrective action to avoid a collision. A similar result can also be achieved 10 if the remote vehicle were connected and received data from the RSE and the sensors that an opposing vehicle was attempting a left turn with a predicted collision.

VBSMs also can be used in lane change maneuvers. Such maneuvers can be dangerous if the vehicle changing lanes does not perform the necessary steps to check the safety of the maneuver, e.g., check back and side mirrors and the blind spot. new advanced driver assistance systems, 15 such as blind spot warnings using onboard ultrasound sensors for instance, have been developed to help prevent vehicles from performing dangerous lane changes. However, these systems can have shortcomings when the sensors are dirty or have an obstructed field of view. And existing systems do not try to warn the endangered vehicle of another vehicle attempting a lane change. V2X communication helps solve this issue through applications such as BS/LCW using BSMs, 20 however the vehicle attempting a lane change may be in unconnected vehicle and therefore not able to communicate its intent. VBSMs can help achieve that goal. Similar to the LTA use case, external sensors installed on the surrounding infrastructure can detect and track an unconnected vehicle attempting a lane change maneuver, collect basic information such as speed, acceleration, heading and past trajectory and transmit them to the RSE. The RSE will in turn 25 build the predicted trajectory for the vehicle changing lanes using rule-based and machine learning algorithms, populate the required fields for the VBSM, and broadcast it on the behalf of the unconnected remote vehicle. The endangered vehicle's OBE will then receive the VBSM with information about a vehicle about to merge into the same lane, process it and determine whether the maneuver presents a potential danger and if it should display a lane change warning 30 to the vehicle's driver. If the vehicle changing lanes is a connected vehicle, its OBE can similarly

- receive VBSMs from the RSE about a vehicle in its blind spot and determine whether the lane change maneuver presents a potential danger to surrounding traffic and if it should display a blind spot warning to the vehicle's driver. If both vehicles are connected, both vehicles will be able to broadcast BSMs to each other and enable BS/LCW applications. However, these
- 5 applications will still benefit from applying the same rule-based or machine learning algorithms (or both) on the BSM data as mentioned above to predict, early on, the intent of a vehicle changing lanes with OBEs deciding whether to display a warning or not.

Autonomous vehicles

- The connectivity that is missing in non-connected road entities affects autonomous vehicles.
- 10 Sensors on autonomous vehicles are either short range or have a narrow field of view. They are unable to detect a vehicle, for example, coming around a building on the corner of the street. They are also unable to detect a vehicle that may be hidden behind a delivery truck. These hidden vehicles, if they are non-connected entities, are invisible to the autonomous vehicle. These situations affect the ability of autonomous vehicle technology to achieve a level of safety
- 15 required for mass adoption of the technology. A smart intersection can help to alleviate this gap and aid acceptance of autonomous vehicles by the public. An autonomous vehicle is only as good as the sensors it has. An intersection equipped with a smart RSE, can extend the reach of the onboard sensors around a blind corner or beyond a large truck. Such an extension will allow autonomous and other connected entities to co-exist with traditional non-connected vehicles.
- 20 Such coexistence can accelerate the adoption of autonomous vehicles and the advantages that they bring.

The virtual connected ground transportation environment includes VBSM messages enabling the implementation of vehicle to vehicle (V2V), vehicle to pedestrian (V2P), and vehicle to devices (V2D) applications that would have been otherwise difficult to implement.

- 25 The system can use machine learning to quickly and accurately generate the fields of data required for the various safety messages, pack them into a VBSM message structure and send the message to ground transportation entities in the vicinity, using various media, such as, but not limited to, DSRC, WiFi, cellular, or traditional road signs.

Virtual personal safety messages (VPMS)

The ground transportation environment can encompass not only non-connected vehicles but also non-connected people and other vulnerable road users.

In some regulatory regimes, connected vulnerable ground transportation entities would continuously send personal safety messages (PSMs). PSMs contain, among other information, the location, heading, speed, and future path of the vulnerable ground transportation entity. Connected vehicles and infrastructure can receive these messages and use them to create a map that includes the vulnerable entities and enhances the level of safety on the ground transportation network.

Therefore, the virtual connected ground transportation environment can bridge the gap between the future when most vulnerable ground transportation entities are expected to be connected and the current time when most vulnerable ground transportation entities have no connectivity. In the virtual connected ground transportation environment, smart traffic lights and other infrastructure installations can use sensors to track all vulnerable ground transportation entities (connected, non-connected) and generate VPSMs on their behalf.

A VPSM message can be considered a subset of a PSM. The VPSM need not contain all fields required to create a PSM but can contain data needed for safety assessment and prevention of dangerous situations and can include localization information including location, heading, speed, and trajectory. In some cases, nonstandard PSM fields may also be included in a VPSM, such as intent, posture, or direction of look of a driver.

The system can use machine learning to quickly and accurately generate these fields, pack them into a VPSM message structure, and send it to ground transportation entities in the vicinity using various media, such as, but not limited to, DSRC, WiFi, cellular, or traditional road signs.

VPSM messages enable the implementation of pedestrian to vehicle (P2V), pedestrian to infrastructure (P2I), pedestrian to devices (P2D), vehicle to pedestrian (V2P), infrastructure to pedestrians (I2P), and devices to pedestrians (D2P) applications that would have been otherwise difficult to implement.

Figure 16 depicts a pedestrian 8102 crossing a crosswalk 8103. The crosswalk 8103 can be at an intersection or a mid-block crosswalk across a stretch of road between intersections. A camera 8101 is used to monitor the sidewalk 8104. The global locations of the boundaries of the field of

view 8105 of the camera 8101 can be determined at the time of installation. The field of view 8105 is covered by a predetermined number of pixels that is reflected by the specifications of camera 8101. A road entity 8102 can be detected within the field of view of the camera and its global location can be calculated. The speed and heading of the road entity 8102 can also be
5 determined from its displacement at s times. The path of the road entity 8102 can be represented by breadcrumbs 8106 which is a train of locations that the entity 8102 has traversed. This data can be used to build a virtual PSM message. The PSM message can then be broadcast to all entities near the intersection.

Traffic Enforcement at Non-Signalized Intersections and Behavioral Enforcement

10 Another useful application of the system is traffic enforcement at non-signalized intersections (e.g. stop sign, yield sign) and enforcement of good driving behavior anywhere on the ground transportation network.

As a byproduct of generating VBSMs and VPSMs, the system can track and detect road users who do not abide by traffic laws and who are raising the probability of dangerous situations and
15 collisions. The prediction of a dangerous situation can be extended to include enforcement. Dangerous situations need not end in collisions. Near misses are common and can raise the stress level of drivers leading to a subsequent accident. The frequency of near misses is positively correlated with the lack of enforcement.

Additionally, using VBSMs the system can detect improper driving behaviors such as abrupt
20 lane changes and other forms of reckless driving. The data collected by the sensors can be used to train and enable machine learning models to flag ground transportation entities engaging in dangerous driving behaviors.

Enforcement authorities usually enforce the rules of the roads for ground transportation entities including vulnerable road users, but the authorities need to be present in the vicinity of the
25 intersection to monitor, detect, and report violations. By tracking non-connected ground transport entities including vulnerable road users using VBSMs and VPSMs, smart RSEs could play the role of enforcement authorities and enforce the rules of the roads at intersections. For example, a non-connected vehicle tracked by a smart RSE could be detected to violate a stop or yield sign, could be identified, and could be reported to authorities. Similarly, a vulnerable road user near an

intersection tracked by a smart RSE could be detected to unlawfully cross the intersection, could be identified, and could be reported to authorities.

For enforcement and other purposes, ground transportation entities may be identified using unique identification including but not limited to plate number recognition. Vulnerable road
5 users may be identified using biometric recognition including but not limited to facial, retina, and voice wave identifications. In special cases that include civil or criminal investigations, social media networks (e.g., Facebook, Instagram, Twitter) may be also used to support the identification of a violating ground transportation entity or vulnerable road user. An example of leveraging social networks is to upload captured pictures of the violator on the social network
10 and request users of the social network who recognize the violator to provide enforcement authorities with intelligence that will help identify the violator.

Other implementations are also within the scope of the following claims.

Claims

1. A method comprising

using electronic sensors located in a vicinity of a crosswalk that crosses a road to monitor an area in or nearby the crosswalk, the electronic sensors generating motion data about
5 vulnerable roadway users who are in and nearby the crosswalk, the motion data including location and direction of the vulnerable roadway users,

applying the generated motion data to a machine learning model running in equipment located in the vicinity of the crosswalk to make a prediction that one of the vulnerable roadway users is about to enter the road in or nearby the crosswalk, the prediction being made prior to the
10 vulnerable roadway user entering the road in or nearby the crosswalk, the machine learning model having been trained using motion data generated in the vicinity of the crosswalk and indicative of intent or behavior of vulnerable roadway users who previously had been in or nearby the crosswalk, the motion data comprising location, speed, acceleration, and orientation, and

15 before the vulnerable roadway user has entered the road in or nearby the crosswalk, transmitting a warning to at least one of: a device associated with the vulnerable roadway user, a device associated with another ground transportation entity approaching the crosswalk on the road, or a road sign configured to alert the vulnerable roadway user or a driver.

2. The method of claim 1 in which the vulnerable roadway user comprises a pedestrian,
20 animal, or cyclist.

3. The method of claim 1 in which the device associated with the vulnerable roadway user comprises a smart watch or other wearable, a smart phone, or another mobile device.

4. The method of claim 1 in which the other ground transportation entity comprises a motorized vehicle.

25 5. The method of claim 1 in which the device associated with the other ground transportation entity comprises a smart phone or another mobile device.

6. The method of claim 1 in which the machine learning model is provided to the equipment located in the vicinity of the crosswalk by a remote server through the Internet.

7. The method of claim 1 in which the machine learning model is generated at the equipment located in the vicinity of the crosswalk.
8. The method of claim 1 comprising training the machine learning model using motion data generated by the sensors located in the vicinity of the crosswalk.
- 5 9. The method of claim 1 comprising sending motion data generated by the sensors located in the vicinity of the crosswalk to a server for use in training the machine learning model.
10. The method of claim 1 comprising segmenting the motion data generated by the sensors located in the vicinity of the crosswalk based on corresponding zones in the vicinity of the crosswalk.
- 10 11. The method of claim 1 comprising using the electronic sensors to generate motion related data representing physical properties of the vulnerable road user.
12. The method of claim 1 comprising deriving trajectory information about the vulnerable road user from motion data generated by the sensor.
13. An apparatus comprising
- 15 equipment located at a level crossing of a transportation network, the level crossing comprising an intersection of a road, a pedestrian crossing, and a rail line, the equipment comprising
- 20 inputs to receive data from sensors oriented to monitor road vehicles and pedestrians at and near the level crossing and to receive phase and timing data for signals on the road and on the rail line,
- a wireless communication device to send to a device of one of the ground transportation entities, pedestrians, or rail vehicles on the rail line, a warning about a dangerous situation at or near the level crossing,
- a processor, and
- 25 a storage for instructions executable by the processor to
- store a machine learning model that can predict behavior of ground transportation entities at or near the level crossing at a current time, the machine learning model

the basic safety messages, virtual basic safety messages, personal safety messages, and virtual personal safety messages comprising location, heading, and speed information about the road vehicles or pedestrians,

5 the basic safety message and the personal safety message being received from the road vehicles or pedestrians,

the virtual basic safety messages and the virtual personal safety messages comprising location, heading, and speed information about the road vehicles and the pedestrians reconstructed from the data received from the infrastructure sensors,

10 applying the received data from the infrastructure sensors and the data from the basic safety messages, the virtual basic safety messages, the personal safety messages, and the virtual personal safety messages to a machine learning model trained to identify dangerous driving or walking behavior of one of the road vehicles or pedestrians, and

automatically reporting the dangerous driving or walking behavior to authorities.

21. The method of claim 20 comprising
15 identifying the road vehicles based on plate number recognition.

22. The method of claim 20 comprising
identifying the pedestrians based on biometric recognition.

23. The method of claim 20 comprising
identifying the road vehicles or pedestrians based on social networking.