



Implementing the RSS Model on NHTSA Pre-Crash Scenarios

Table of Contents

WHAT IS RSS AND WHY IS IT NECESSARY?	3
DEFINITIONS OF A SAFE DISTANCE.....	4
DEFINING DANGEROUS SITUATIONS.....	10
PEDESTRIANS AND LIMITED VISIBILITY	13
RESPONSIBILITY ASSIGNMENT - EXAMPLE	16
IMPLEMENTING RSS ON NHTSA PRE-CRASH SCENARIOS	17
• ONE-WAY TRAFFIC SCENARIOS	17
• CUT-IN AND DRIFTING SCENARIOS	18
• TWO-WAY TRAFFIC SCENARIOS	19
• MULTIPLE GEOMETRY AND RIGHT-OF-WAY SCENARIOS	20
• VULNERABLE ROAD USERS AND OCCLUSION SCENARIOS.....	23
CONCLUSION	24

Social acceptance of autonomous vehicles relies, first and foremost, on the ability of technology developers to measure risks and above all to ensure safety. As of today, the concepts of “dangerous situations”, “proper response” and the notion of “blame” remain a subjective matter left for interpretation. Ensuring safety by showing that we have tested the most miles, or that we use the “best practice” is simply not enough. Even programming AVs to follow the letter of the law doesn’t ensure low risk. True safety assurance, alongside societal acceptance of AVs, is a challenge the industry must overcome together. That is where the **“Responsibility-Sensitive Safety” (RSS)** model comes in, as one of the underlying objectives of RSS is to evoke discussion among industry groups, car manufacturers, regulatory bodies and the general public. Now is the time for the industry to take the debate over AV safety to the next level. RSS can be a useful starting point.

What is RSS and Why is it Necessary?

At its core, the RSS model is designed to formalize and contextualize human judgment regarding all driving situations and dilemmas. Once we do this, an AV can be programmed to follow these definitions of safety, thereby making it possible for AVs to share the road with human drivers for the next couple of decades. This includes notions of safe distance and safe gaps when merging and cutting in, right-of-way, how to define safe driving with limited sensing and visibility, and more. This formalization is then translated into a set of rules in four essential realms. First, RSS defines **Safe Distances** for all driving scenarios, from one-way traffic to junctions and multiple geometry scenarios. Secondly, RSS defines exactly what is considered to be a **Dangerous Situation**, as a derivative of all the semantics and rules defined in the “human judgment” formalization. The third essential definition in the model is the **Proper Response** that needs to be taken in order to evade a dangerous situation. By formalizing those first three realms, we create a set of parameters that are lacking in today’s development of autonomous vehicle safety systems. The absence of those parameters is an issue that must be addressed as part of the effort to make AVs safe and more “human-like”. The human perception of safety consists of a subjective and interpretable set of rules and principles. In order for the AV to comply with the human perception of safety, we must couch it in mathematical formulas such that it would be possible to program machines accordingly.

To better understand this motivation let us use a few simple examples of what exactly needs to be defined. Human drivers all have a common understanding that if you hit someone from behind, you are to blame. However, not only is this rule of thumb not always true, but we also do not know with precision what it means to keep a safe distance from the vehicle ahead in complex configurations. Another area where human driving norms lacks clarity is the definition of a safe cut-in/merge. Although the average driver probably has a notion of what seems like a safe cut-in, this would only be a subjective and interpretable notion, as there is no precise and measurable definition of what constitutes a safe cut-in. In other words, **we cannot precisely tell when we are in a dangerous situation, we simply rely on our common sense and subjective judgment. As opposed to humans, machines require precise definitions which are based on precise measurements rather than subjective and ambiguous definitions.** Therefore, implementing the perception of the safety of common situations into AV software would be impossible without establishing a decision-making mechanism that is based on definitions more precise than common sense, norms, and traffic law.

Once we have defined what constitutes a dangerous situation, what caused it, and how to respond to it, we can add on the fourth layer to RSS: the definition of **Responsibility**. By formally defining the parameters of the dangerous situation and proper response, we can say that responsibility is assigned to the party who did not comply with the proper response. **Therefore, the RSS model guarantees that when applying it to any “driving policy” (the decision-making mechanism of the AV), the self-driving car will never initiate a dangerous situation and thus, it will never cause an accident.** Eventually, we believe that building a foundation that formalizes all aspects of human judgment in the context of driving will facilitate creation of **a formal “seal of safety” for autonomous vehicles** that will allow car manufacturers to limit the inherent risk in deploying AVs on public roads alongside human drivers to extremely low levels. We believe these AVs would be responsible for accidents at a rate 1000x better than human-driven vehicles.

The RSS model was first introduced as an academic paper¹ by Professor Amnon Shashua, Professor Shai Shalev-Shwartz, and Shaked Shammah. Since it was published, it was presented at several [conferences, blogs, short videos](#), and in other manners. **In this article, we aim to go beyond the conceptual and demonstrate the RSS methodology using real-world pre-crash scenarios in order to demonstrate the soundness of the model. The goal is to show that when the model determines that an agent responded properly to any of a vast variety of pre-crash scenarios, a human in the same situation would have reached the same conclusion.**

The following analysis is based on research done by the National Highway Traffic Safety Administration (NHTSA)², which defines, and statistically describes, a typology for pre-crash scenarios for light vehicles based on the 2004 General Estimates System (GES) crash database. The research presents 37 pre-crash scenarios that represent about 99.4 percent of all light-vehicle crashes. We aim to demonstrate the soundness of RSS by applying it to those 37 scenarios.

Definitions of a Safe Distance

Before diving into the scenario analysis, we will first outline the essential definition of the RSS model, to facilitate an adequate case inquiry. We begin with defining safe longitudinal and lateral distances in various situations. First, we define **Safe Longitudinal Distance in One-Way Traffic**, as depicted in Figure 1. Generally speaking, **the longitudinal distance between the blue car followed by the red car is considered safe if the red car will avoid a collision even if the blue car abruptly applies full braking force.** Specifically, the distance is considered safe if an accident is not possible, even if the blue car applies full braking force ($a_{max,brake}$) and during its response time (ρ) the red car accelerates at maximum acceleration ($a_{max,accel}$) and then immediately brakes by at least the minimum reasonable braking force ($a_{min,brake}$) that is likely to be used by a human driver in this situation. In order to simplify definitions in this article, from this point forward, we will refer to this as the **“stated braking pattern”**³.

¹ [On a Formal Model of Safe and Scalable Self-driving Cars](#)

² [Pre-Crash Scenario Typology for Crash Avoidance Research](#) (NHTSA), 2007

³ Applying maximum acceleration ($a_{max,accel}$) during the response time (ρ) and then immediately braking by at least the minimum reasonable braking force ($a_{min,brake}$).

The model demonstrates a conservative approach by assuming the “worst case scenario”, meaning that the red car accelerates to the maximum speed during its response time and then does not apply the full braking force when needed. The value of a worst-case scenario is that it removes the need to estimate “road user intentions” and provides a guarantee that applies regardless of the action of other road users. In order to simplify the definitions later used in the article, we assume that whenever our agent needs to apply full brakes, it will follow this *braking pattern*. **The reasonable value for all parameters defining safe longitudinal distance**

($a_{max,brake}$, $a_{max,brake}$, $a_{max,accel}$, ρ), along with all other parameters defined in the course of this article, should be determined through collaborative discussions amongst the the industry and automotive standards-bodies, likely facilitated by regulators.

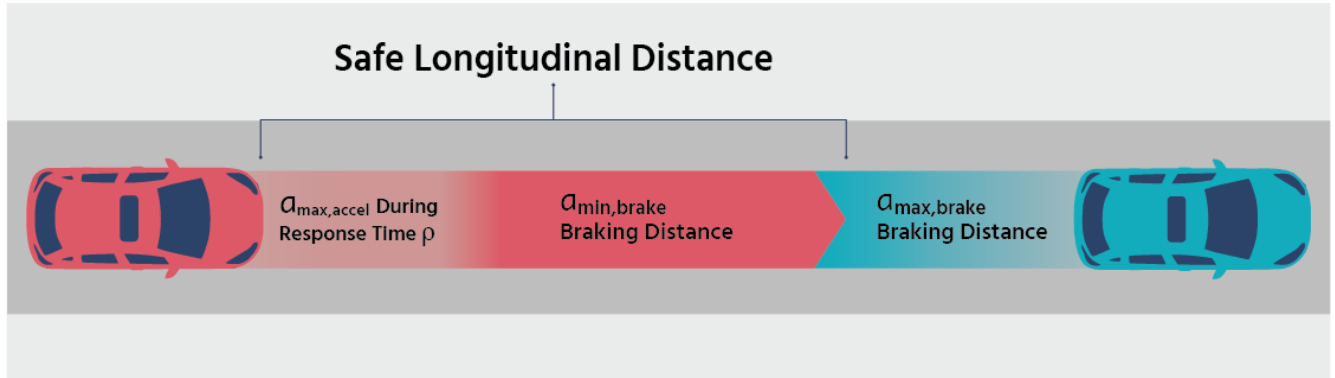


Figure 1: Red car's safe longitudinal distance

Next, we expand the definition of **Safe Longitudinal Distance to Two-Way-Traffic** scenarios. As depicted in Figure 2, the red car is driving towards the blue car. The blue car is driving in its lane while the red car encroaches into the opposite lane (for overtaking, avoiding an obstacle, etc.). The longitudinal distance between both cars is considered to be safe when a collision is not possible even if the red car conducts the Stated Braking Pattern³ until coming to a full stop while the car driving in its lane (i.e. the blue car) brakes to a full stop with at least the minimal reasonable braking force ($a_{min, brake, correct}$). We set $a_{min, brake, correct}$ to be the reasonable braking force expected from a car driving in its correct lane and direction, where $a_{min, brake, correct}$ is smaller than $a_{min, brake}$. In other words, we expect the car that is driving in its lane to brake more moderately than the car driving in the wrong direction.

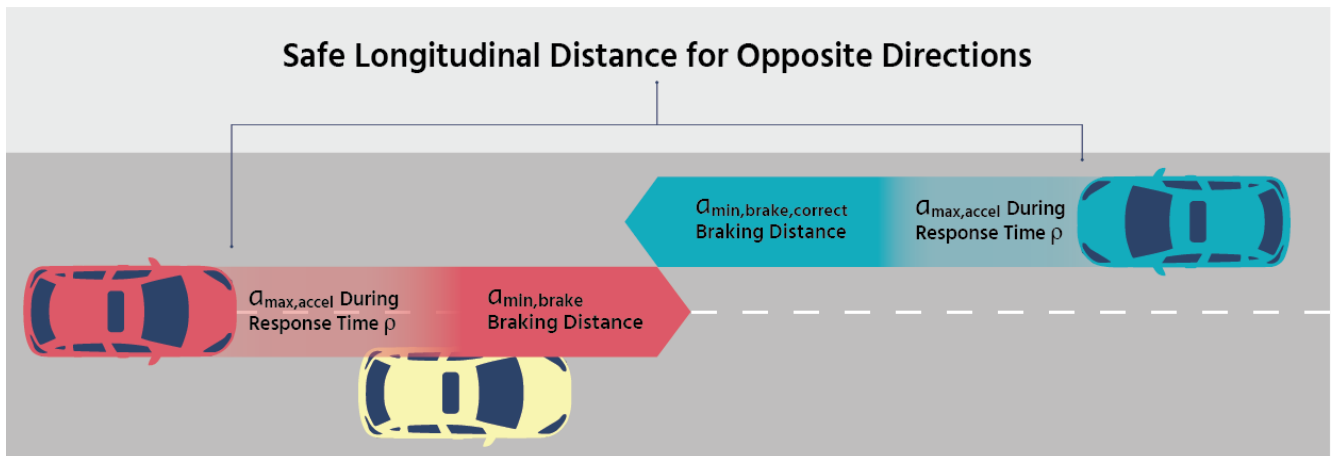


Figure 2: The red car is overtaking the yellow car while encroaching the blue car's lane. The red car is obligated to brake more strongly than the blue car, as the blue car is the one who is driving in the correct lane

Roughly speaking, for a collision to occur between two vehicles, they must be close to each other both longitudinally and laterally. For the longitudinal axis, we have already formalized the notion of "being close" using the Safe Distance formula. We will now do the same for **Lateral Safe Distance**. The lateral distance between two cars is considered safe if both cars apply the Stated Braking Pattern³ until they reach zero lateral velocity without colliding. This definition also takes into account possible fluctuations that are typical to lateral maneuvers by adding some notion of margin, as depicted in Figure 3.




Figure 3: The lateral safe distance between two cars

We now turn to deal with scenarios in which there are **multiple route geometries** that overlap in a certain area. Examples include roundabouts, junctions, and highway merge situations. In many cases, one route has priority over others, and vehicles riding on it have the right-of-way. Moreover, such road geometries often do not share a coherent longitudinal and lateral coordinate system. When facing scenarios of multiple route geometries, the definitions we described thus far provide only a limited solution. Therefore, expansions of the model and adjustments to the definitions of safe distances are necessary.

Multiple route geometry is intimately related to the rule of “right-of-way”. The common sense concept that “right of way is given rather than taken” must be formalized in order for the definition of safe distances (longitudinal and lateral) to correctly reflect human judgement. In other words, the blame for a collision does not always reside with the agent who did not obey the right-of-way. For example, human drivers prefer to be safe rather than correct, and so will yield their right-of-way when necessary to avoid an accident. The following definitions are a formalism of this common-sense principle as well as other route priority rules.

First we depict a definition that will later help us deal with right-of-way and route priority: **Longitudinal Ordering for Two Routes of Different Geometries**. Consider two vehicles, blue and red, driving on different routes that intersect, as depicted in Figure 4. RSS determines that, in the case of Figure 4(a), the red vehicle is longitudinally in front of the blue vehicle since it is closer to the intersection of the two routes. This definition helps us to determine the proximity of the vehicle to the critical point of route intersection with no consideration of route priority. Based on this definition, we will later formalize the principle that the right-of-way is given and not taken.




Figure 4(a): red car is in front of blue car - $D_{red} < D_{blue}$.




Figure 4(b): blue car is in front of red car - $D_{blue} < D_{red}$.

In the same manner that we formalized previous safe distance definitions, we define the **Safe Longitudinal Distance for Multiple Geometries**. First we formalize the simplest notion of route priority: a vehicle which has to give way should be able to stop safely before encroaching the prioritized lane. In Figure 5 the red car has priority, and the blue car needs to keep distance such that after applying the Stated Braking Pattern³, it will still not encroach upon the route of the red car.




Figure 5: the distance in the figure to the left is considered safe since the blue car can stop safely without encroaching the red car's pathway

Consider a situation in which the non-prioritized vehicle cannot stop safely, such that the encroachment into the prioritized lane is inevitable. This is where the formalization of the notion that "right-of-way is given and not taken" is needed. As depicted in Figure 6, when the blue car (non-prioritized) is in front of the red car, the longitudinal distance is considered safe if the longitudinal ordering is maintained after the blue car applies maximum braking force until reaching a zero velocity, while the red car conducts the Stated Braking Pattern³. This definition will later allow us to compel prioritized cars to take action to prevent imminent collisions, as we will later define in the Proper Response to this situation.




Figure 6: the distance in the figure to the left is considered safe since the red car is able to stop safely if the blue car is in front of it and encroaching to its lane

Now we depict a derivative of the definition above, in order to cope with situations in which the encroachment of the non-prioritized car into the prioritized lane is inevitable, but occurs at speeds such that the prioritized car can avoid a collision by maneuvering laterally (e.g. changing lane, passing the other car, evasive maneuver). We define the time interval $[t_{red,min}, t_{red,max}]$ and $[t_{green,min}, t_{green,max}]$ to be all times in which the red car and the blue car respectively, can arrive to the intersection point of the two routes if each of them apply longitudinal accelerations in the range $[a_{max,brake}, a_{max,accel}]$ during the response time, and then apply longitudinal braking in the range $[a_{min,brake}, a_{max,brake}]$ until reaching zero velocity. The distance in this situation is considered safe as long as those two time intervals do not intersect. Basically, those definitions of time intervals and velocities are applicable for situations such as congested junctions, a car slowly sliding into a junction, etc. See the illustration in Figure 7.




Figure 7(a): time intervals overlap. Thus, the distance is considered non-safe




Figure 7(b): due to lane changing of the red car the time intervals no longer overlap. Thus, now the distance is safe

Defining Dangerous Situations

Based on the basic definitions of keeping a safe distance, we can now define what constitutes a **Dangerous Situation**. RSS determines that a situation is dangerous for two cars if *both* the longitudinal and lateral distances between them are non-safe. For example, see Figure 8.




Figure 8(a): non dangerous-
both longitudinal and lateral
distance is safe

Figure 8(b): non
dangerous- only lateral
distance is non-safe

Figure 8(c): non dangerous-
only longitudinal distance is
non-safe

Figure 8(d): dangerous- both
longitudinal and lateral
distance is non-safe

Next we define the “Proper Response” to dangerous situations. The Proper Response is a set of actions that the car must take in order to avoid a collision. For two cars driving in the same direction as depicted in Figure 9, the Proper Response should be as follows:

- In case the front car applies full brakes until reaching a full stop, the Proper Response for the rear car is to conduct the Stated Braking Pattern³ until reaching a full stop.
- For two cars driving in opposite directions as depicted in Figure 10 where the red car is driving in the wrong direction, namely not in its lane, both cars must comply with a proper response. The red car is obligated to conduct the Stated Braking Pattern³ until full stop, while the car driving in its lane (i.e. the blue car) is compelled to brake to full stop with at least the minimal reasonable braking force ($a_{min,brake,correct}$).

- In case the lateral distance is not safe, both should brake laterally by at least $a_{min,brake}^{lat}$ until reaching a safe lateral distance.
- **Minimal evasive effort** deals with cases in which extra caution is applied to prevent potential situations in which responsibility might be shared. For example, in Figure 8(c), say the blue car performs a cut-in to the red car's lane at a non-safe longitudinal distance, while the red car was behaving properly. In this case, the blue car did not behave correctly. Nevertheless, we expect the red car to make an effort to evade a potential collision, by applying moderate braking force in order to evade the Dangerous Situation. **This ensures that Dangerous Situations cannot last long, even if our agent wasn't the one who caused it.**

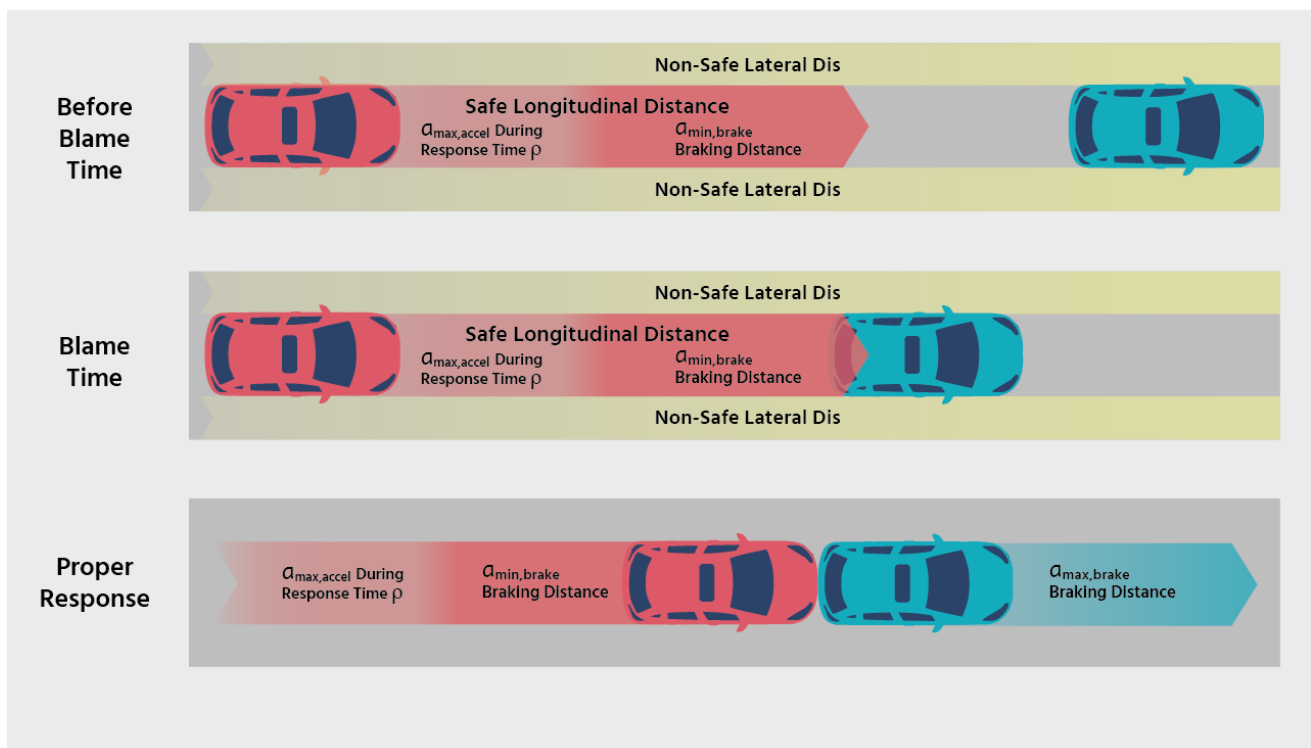


Figure 9: The proper response for the red car when following the blue car in the same direction

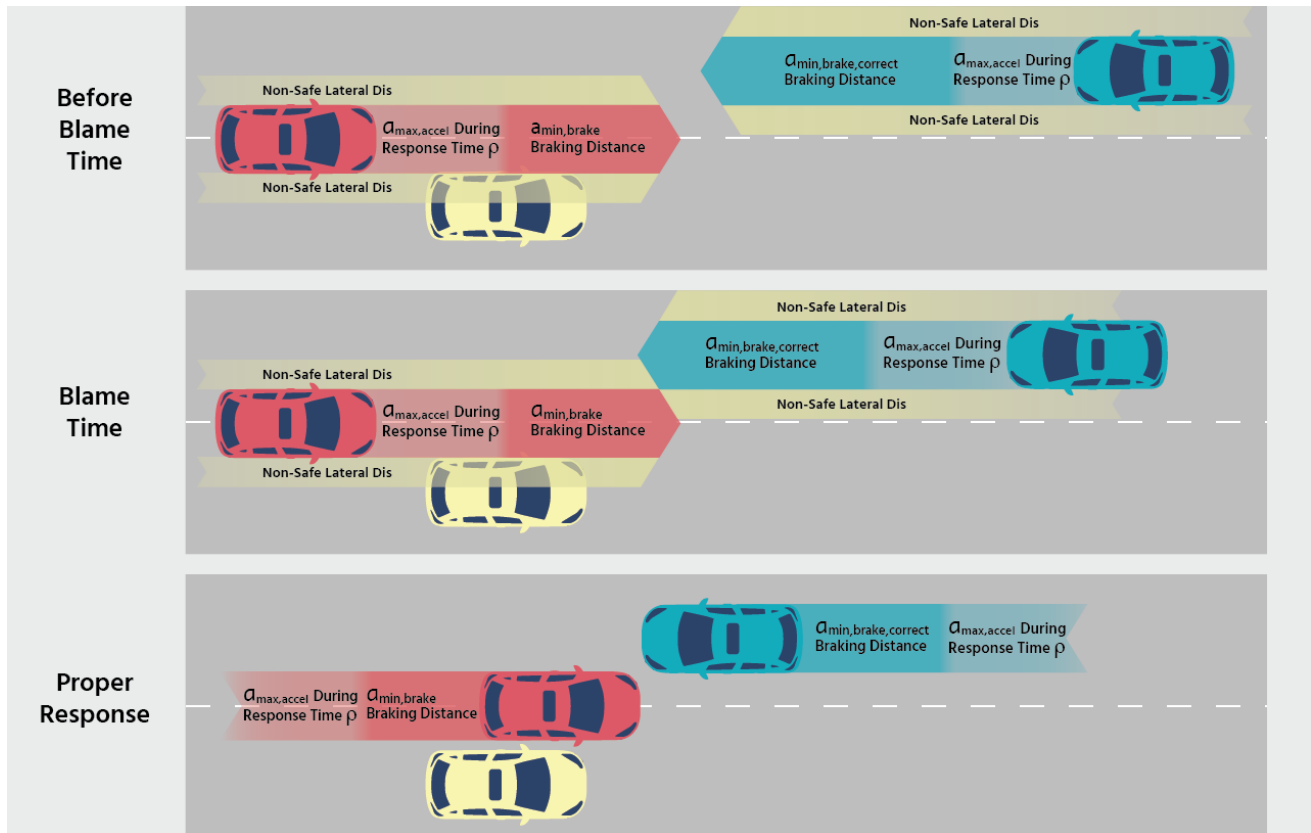


Figure 10: The proper response for the red car when driving towards the blue car in opposite directions

Lastly, we define the **Proper Response to a Dangerous Situation** and **Responsibility for Two Routes of Different Geometry**, in the same manner we did for one-way and two-way traffic. The actions derive from the situation before the **Blame Time**, i.e. the time in which the Dangerous Situation was triggered:

- If before the blame time the longitudinal distance was safe as in Figure 5, the vehicle on the prioritized route can drive normally. The non-prioritized car must conduct the Stated Braking Pattern³ until reaching a full stop.
- If the non-prioritized vehicle is in *front* of the prioritized vehicle, and cannot stop safely (as in Figure 6), the vehicle in front can drive normally. The rear car must then conduct the Stated Braking Pattern³ until reaching a full stop. **This definition of Proper Response is the formalization of the principle “right-of-way is given and not taken”, as it compels even the car on the prioritized route to take action to prevent an imminent collision.**

Note: One may argue that according to this definition, there could be a situation in which RSS assigns blame only to the prioritized car, while the non-prioritized car holds no responsibility, as it has no constraints in the proper response definition. It is true that human judgment would often assign blame to some extent, also to the car who violated the right-of-way, even if the prioritized car could have avoided the accident. This is a good opportunity to look at the situation from an AV perspective and in light of the RSS definition of

soundness - **the RSS model will never clear the blame from our agent (i.e. the AV) where human judgment states that it is to blame.** Namely, if the AV is on the prioritized lane, it will never take right-of-way unless it is given to it, whereas if it is in the non-prioritized lane, it will be able to merge into traffic in an assertive, yet safe, manner. **This reflects the conservative approach of RSS regarding blame assignment in combination with an assertive and “human-like” driving style, with no need to trade-off one for the other.**

- To gain comprehensive coverage of all possible situations, the model also copes with cases where neither of the above situations holds. These cases can occur when the vehicle that is in front is slowly encroaching upon the prioritized lane (e.g. a car stopped at a stop sign and then slowly sliding into the junction). In those cases a combination of lateral and longitudinal maneuvers is required to avoid a collision, such as changing lanes or overtaking. In the event that lateral maneuvers are not possible, both cars should brake laterally and longitudinally (each one w.r.t. its own route) according to the Stated Braking Pattern³ in order to avoid a collision. This definition compels both cars, even the prioritized car, to take action to avoid a collision.

Pedestrians and Limited Visibility

We now turn to Vulnerable Road Users, or VRUs (pedestrians, cyclists, animals, etc.) and occlusion-related accidents (i.e. accidents involving fully or partially hidden objects). Determining responsibility in collisions involving pedestrians (or other road users) follow the same ideas described in previous subsections. Even though the concepts are similar, an adjustment to the parameters that constitute the definitions of Safe Distance and Proper Response is needed. First, we will delineate the methodology that the RSS model exercises to cope with pedestrian-related scenarios. The second piece of the puzzle relates to obstructions and limited sensing scenarios. Accidents related to pedestrians and other VRUs often take place in an obstructed-view environment and under conditions of limited visibility (i.e. limited sensing, in the context of an AV). The RSS model ties those two spheres together.

For pedestrians, the first step would be to specify pedestrian routes and their priority. In some cases, a pedestrian's route is well defined (i.e. marked crosswalk or a sidewalk on a fast road). In other cases like a typical residential street when the pedestrian's intended route is often unclear, we follow an approach that considers all possible trajectories. We constrain the change of heading and assume that in an emergency, after the response time, the pedestrian will continue in a straight line. If the pedestrian is standing, we contemplate all possible lines originating from their current position. The priority is set according to the type of the road and possibly based on traffic lights. For example, in a typical residential street, a pedestrian has the priority over the vehicles, therefore vehicles must yield and be cautious concerning pedestrians. In contrast, there are roads with a sidewalk where the common sense judgment is that vehicles should not be worried that a pedestrian on the sidewalk will suddenly start running into the road. Cars have the priority in this case. Another example is a zebra crossing with a traffic light, where the priority is set dynamically according to the light. Of course, priority is given not taken, so even when pedestrians do not have priority, cars must brake and let them pass, assuming they entered the road at a

safe distance.

Let us illustrate the idea with three examples. The first example is a pedestrian that stands on a residential road. All possible routes originating from their current position are assigned to the pedestrian. The safe longitudinal distance with respect to each of these virtual routes is quite short. For example, a pedestrian's safe longitudinal distance is 50cm assuming a delay of 500ms and maximal reasonable acceleration and braking⁴ of $2 \frac{m}{s^2}$. Therefore, the vehicle must be in a kinematic state such that applying the proper response would allow it to remain outside a 50cm radius around the pedestrian.

Consider a second example where a pedestrian stands on the sidewalk in front of a zebra crossing, facing a red light while a vehicle approaches the zebra crossing at a green light. Here, the vehicle route has the priority and the vehicle can assume that the pedestrian will stay on the sidewalk. If the pedestrian enters the road while the vehicle is at a safe longitudinal distance (from the vehicle's point of view), then the vehicle must brake ("right-of-way is given, not taken"). However, if the pedestrian enters the road while the vehicle is not at a safe longitudinal distance such that even Stated Braking Pattern³ cannot prevent a collision, then the vehicle is not responsible.

The third example is a pedestrian that runs on a residential road at 10 km per hour (~ 2.7m/s). Using the same parameters as in the first example, the pedestrian's safe distance is roughly 15m. It follows that cars should not enter the pedestrian's route at a distance smaller than 15m. But, if the car entered the pedestrian's route at a distance larger than 15m, the pedestrian did not stop and crashed into the car then it is the pedestrian's responsibility.

In the case of pedestrians who are occluded from the vehicle (i.e. view is obstructed), additional caution is warranted. **The vehicle should adjust its speed such that if a child would emerge from behind some object, then there would be no accident.** By assuming the maximal speed of the child (or another pedestrian), to some reasonable amount, the vehicle can adjust speed accordingly to avoid a collision with a pedestrian even in a worst-case scenario where the pedestrian emerges from behind some sensing obstruction (e.g. a parked car) even at that maximal speed.

However, this yields behavior that is too defensive. Consider a typical scenario where a vehicle is driving next to a line of parked cars. A child that intends to run into the road (e.g. chasing a ball) at a speed of 10km/h, might be occluded by one of these parked cars. As illustrated in Figure 11, a pedestrian obstructed behind an obstacle will be revealed to the vehicle's camera only when the longitudinal distance between them is smaller than 0.3m. However, the safe longitudinal distance in the pedestrian's route is around 15m, which is typically more than the width of the residential road. Therefore, the vehicle is already at a non-safe longitudinal distance, and it therefore, must be able to stop before entering the pedestrian's route. However, the braking distance of a car driving at 1m/s (with the reasonable set of parameters) is more than 0.4m. It follows that we should drive slower than 1m/s in this scenario, even if we drive at a lateral distance of, say, 5m, from the parked cars. This is too defensive and does not reflect normal human behavior. What assumptions does the human driver make, which allows driving faster? The first justification a human driver is making relies on the law - driving slower than the speed limit often feels safe. **While AVs will obviously follow the speed limit, we seek a stronger notion of**

⁴ The estimated acceleration of Usain Bolt is $3.09 \frac{m}{s^2}$

safety. As explained below, the RSS model proposes a responsibility definition for this scenario which leads to safer driving without constraining speed to unreasonable levels.




Figure 11: the front facing camera of an AV can first sense an obstructed pedestrian at a distance of 0.3m

In the next subsection, the model deals with occlusions, and in particular with occluded pedestrians. Scenarios of this type are often hard to judge, but under the known constraints, the RSS model can formalize the responsibility structure of those situations with three main definitions. The first definition is the **Exposure Time**. As depicted in Figure 12, the Exposure time is defined as the time in which an object can first be seen by the vehicle (for a human driver- regardless of when he actually saw it).




Figure 12: the exposure time is defined as the first moment in which the pedestrian can first be seen by the vehicle

Secondly, we define responsibility in such cases. Responsibility should be divided in two: **Responsibility Due to Unreasonable Speed** and **Responsibility for Collision with Occluded Pedestrian**. As for the former, a car should be held responsible where before the blame time there was a safe lateral distance, and it was driving unreasonably fast. Namely, its average velocity from the exposure time until the blame time was larger than v_{limit} , where v_{limit} is a parameter associated with posted speed limit and other road conditions. As for the definition of **Responsibility for Collision with Occluded Pedestrian**, RSS determines that a vehicle is not responsible for an accident with an occluded pedestrian if these two conditions hold: first, the vehicle did not accelerate during the response time (the time interval that starts at the exposure time until the end of the response time) and performed a longitudinal brake until the accident or until arriving to a full stop. Second, the average velocity of the vehicle from the exposure time until the collision time must be smaller than the average velocity of the pedestrian at the same time interval. Otherwise, the vehicle is at fault.

The first condition tells us that the velocity of the vehicle is monotonically non-increasing from the exposure time until the collision time, and the second condition tells us that the average velocity of the vehicle is smaller than that of the pedestrian. The responsibility definition tells us that at the collision time, we expect that either the velocity of the car will be significantly smaller than that of the pedestrian, or that both the vehicle and the pedestrian will be traveling at a very slow speed. In the latter case, the damage should not be severe and in the former case, it complies with “common sense” that the pedestrian is to be blamed for the accident.

Responsibility Assignment - Example

At this point, and after covering most of the essential definitions of RSS, let us outline a simple example to understand the motivation of this model better. Consider an accident that occurs from the situation depicted in Figure 8. Two main scenarios can yield an accident in this situation. The first is if the blue car violated the red car's priority and merged into its lane at a non-safe distance. The second is if the blue car did not violate the right-of-

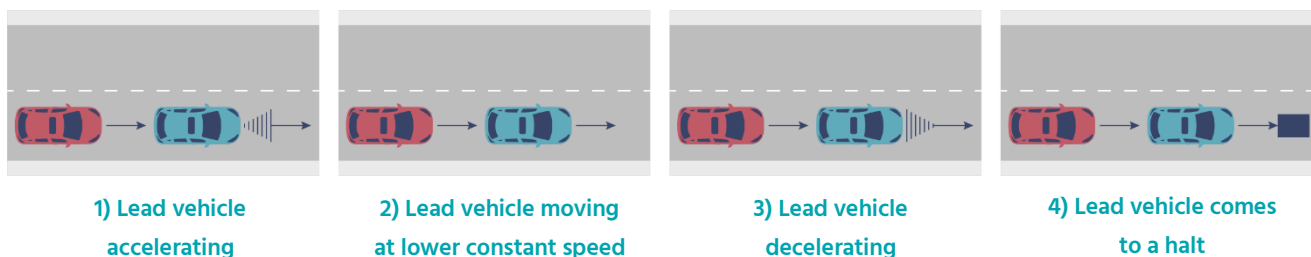
way while merging in at a safe distance, and yet the red car hits the blue car anyway, possibly due to the inattentiveness of the person driving the red car. In many versions of these two scenarios, blame assignment according to human judgement is obvious. For example, if the blue car brutally encroaches the red car's path leaving the red car no time to react. The bigger challenge is in the "grey area", in cases that do not conform to clear-cut categories or common-sense rules. For example, in cases where there is no measurable, clear definition of what it means to give or violate right-of-way. In these cases, by implementing the definitions of dangerous situations and proper responses outlined above, the result is a clear determination of **Responsibility** for the collision. RSS determines that a car should be held responsible for an accident if it did not follow the constraints of the "proper response to a dangerous situation". There are also situations in which more than one agent does not comply with the proper response. The RSS model is also capable of dealing with such situations where responsibility can be assigned to more than one agent.

Implementing RSS on NHTSA Pre-Crash Scenarios

Once we have established a method that allows us to determine who is to blame in all possible road geometry scenarios, we can now examine the soundness of the model by implementing it on representative real-world scenarios taken from the NHTSA research¹. Note that the RSS model copes only with scenarios in a multi-agent environment. Hence, scenarios within the NHTSA research that do not stem from multi-agent scenarios (such as vehicle failure, loss of control, crashing into a stationary object, etc.) are not covered in this article.

One-Way-Traffic Scenarios

Scenarios as depicted in the NHTSA research:




It is quite clear that all of the above scenarios are similar when examining them through the RSS methodology and therefore, we shall consider them as one.


In the above scenarios, dangerous longitudinal situations were initiated when the safe longitudinal distance was first violated. The proper response to this scope of scenarios states that the front car has the liberty to brake up to the maximum braking force while at the same time the rear car must use the Stated Braking Pattern³, as defined by RSS. Thus, assuming an accident occurred as a result of one of these scenarios, the rear car is to be blamed, for not complying with the proper response to this situation.

Cut-In and Drifting Scenarios


Scenarios as depicted in the NHTSA research:



1) Vehicle changing lane-
same direction



2) Vehicle turning left-
same direction



3) Vehicles drifting-
same direction

In scenarios 1 and 2, when a vehicle is changing lanes (i.e. cutting-in), we need to examine whether it was a safe cut-in or not, considering the safe longitudinal distance with regard to the car that is being cut in front of, as depicted in Figure 13:




Figure 13(a): non safe cut-in as the red car is violating
the frontal longitudinal safe distance of the blue car




Figure 13(b): safe cut-in as the red car is not violating
the frontal longitudinal safe distance of the blue car

Assuming a non-safe cut-in, the dangerous situation is initiated when the red car violated both the longitudinal and lateral safe distance of the blue car. Therefore, it is to blame in case of a collision. However, the blue car is expected to initiate a minimal evasive effort in order to evade the dangerous situation that was forced upon it.


In scenario 3 regarding a car drifting from its lane, the longitudinal distance was non-safe to begin with. The situation is not considered dangerous, however, until one of the cars violates the safe lateral distance. Assuming that before the blame time there was a safe lateral distance, then the proper response would be to brake laterally until reaching zero lateral velocity. The car which did not maneuver laterally is already at zero lateral speed, hence only the other car should brake. If that car does not brake laterally then it did not comply with the proper response and thus is responsible. Again, if it would not initiate a separate dangerous situation, we would expect the other car to conduct an evasive effort to avoid the dangerous situation, even though it was caused by the other vehicle.

Two-Way Traffic Scenarios

Scenarios as depicted in the NHTSA research:



1) Vehicle overtaking another vehicle-
opposite direction



2) Vehicle drifting-
opposite direction

In scenario 1, the vehicle is passing another vehicle in the same direction while entering the lane of an oncoming vehicle. In scenario 2, the vehicle is driving straight ahead, without making any deliberate maneuvers but drifting into the lane of the vehicle travelling in the opposite direction. From the RSS standpoint, both scenarios are treated the same.

The dangerous situation is triggered when the vehicles' longitudinal and lateral safe distances first overlap, as depicted in Figure 14 below. RSS determines that prior to the blame time there was a safe lateral distance but a non-safe longitudinal distance. Thus, the proper response would be for the red car to brake laterally and return to a safe lateral distance. Assuming a collision occurs from this situation, the red car is to be blamed for not initiating a proper response.

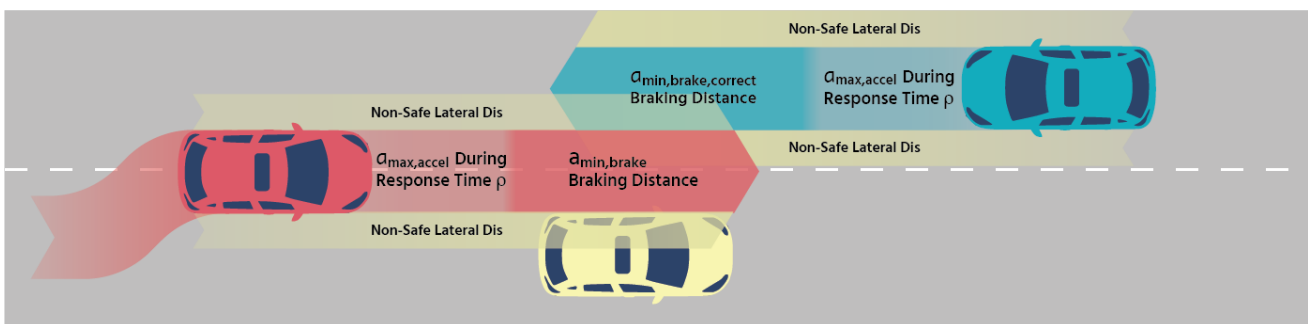


Figure 14: the red car initiated the dangerous situation by maneuvering latterly into the lateral and longitudinal safe distance of the blue car

On the other hand, if an accident stems from the situation depicted in Figure 15, both cars will be held responsible for not initiating the proper response to a dangerous longitudinal distance for vehicles traveling in opposite directions. In other words, because there was a safe longitudinal distance before the blame time, we can determine that the blue car could have avoided the collision if it had complied with the proper response.

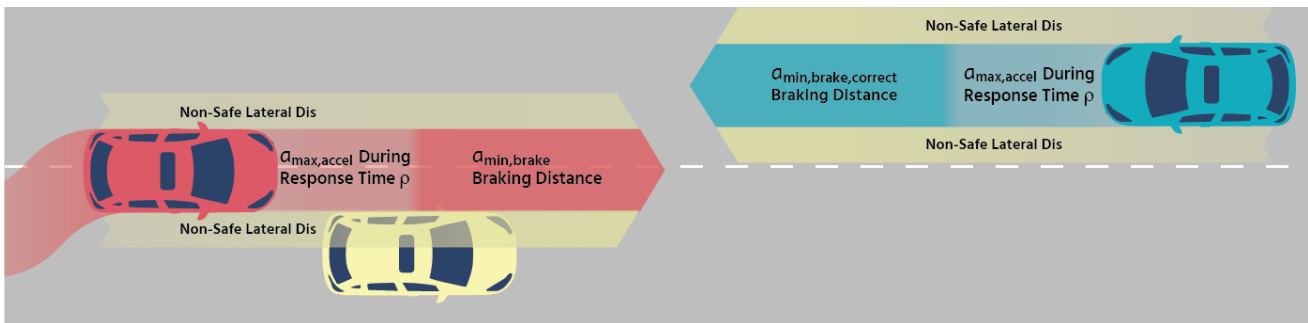



Figure 15: as opposed to the previous scenario, here the lateral maneuver of the red car did not initiate any dangerous situation since it did not encroach to the lateral and longitudinal safe distance of the blue car

Multiple Geometry and Right-of-Way Scenarios

Scenarios as depicted in the NHTSA research:



1) Running a red light

2) Running a stop sign

3) Vehicle stops, but
not giving way

4) Vehicle turning left at
non-signalized junction
against crossing traffic

In the scenarios depicted above, we see a car not respecting the priority assigned to the other car by the traffic light or the stop sign. Dangerous situations in the scenarios above are first initiated when the longitudinal distance, according to the definition of safe longitudinal distance in multiple geometries, is first violated (lateral distance in multiple geometries is essentially always considered non-safe). As mentioned earlier in this article, longitudinal distance in those scenarios is considered safe if one of the three following conditions hold:

1. The non-prioritized car can stop safely before encroaching upon the prioritized lane.
2. If the non-prioritized car cannot stop safely, it implies that the encroachment into the prioritized lane is inevitable. In this case, if the non-prioritized car is in front of the prioritized car (as we previously defined in the "longitudinal ordering" definition), the prioritized car should follow the Stated Braking Pattern³ to

avoid a collision, even though it has priority by law.

3. In cases where neither of the above conditions holds, a combination of lateral and longitudinal maneuvers are required to avoid a collision. In a case where no such maneuver is possible, both cars should brake longitudinally and laterally.

Assuming a collision occurs in one of the situations above, we can clearly disprove the first condition. Namely, the non-prioritized car, intentionally or unintentionally could not stop safely before encroaching upon the prioritized pathway. Now let us examine the case in light of the second condition: as illustrated in Figure 16(a) below, if the non-prioritized car is in front the prioritized car, and the prioritized car can stop safely and avoid a collision, it should definitely do so, as defined in the proper response to this situation.




Figure 16(a): the red car is longitudinally in front of the blue car, whereas the blue car can follow the stated braking pattern and avoid a collision




Figure 16(b): lateral evasive maneuver to pass the red car as it is slowly encroaching the prioritized lane

Where the prioritized car cannot avoid the accident, we shall examine whether the third condition holds. As depicted in Figure 16(b), the lateral evasive maneuver is possible only if the non-prioritized car is travelling at a very low speed or even at a full stop. In cases where no such maneuver is possible, as depicted in Figure 16(c) both cars should comply with the Stated Braking Pattern³ until they reach a full stop. In all the cases above, the agent who did not comply with the proper response is to blame. The above scenario (3) illustrates the necessity of such a definition. In such a case, the non-prioritized vehicle is at a safe longitudinal distance according to condition (1) as long as it is outside the prioritized car's pathway. This condition no longer holds if the non-prioritized car slowly encroaches the prioritized car's pathway. If the encroachment takes place at a moment where the longitudinal distance is safe according to the second condition and a collision occurs, the RSS model determines that the prioritized car is to blame, since it had sufficient time to react. Otherwise, if the encroachment takes place at a

moment where the longitudinal distance is non-safe according to the second condition, this scenario can be judged only by adding in the third condition. Consider a state where before the dangerous situation, both time intervals that we previously defined did not intersect, as depicted in Figure 16(c). In such cases, RSS states that both cars should brake longitudinally and laterally until reaching a full stop. The agent who did not comply with the proper response should be held liable.




Figure 16(c): the two time intervals do not overlap. A dangerous situation will be initiated at the moment of the first overlap

Scenario (4) depicts a situation in which a vehicle is turning left at an intersection and cuts across the path of another vehicle travelling in the opposite direction. If a collision occurs in this situation, it implies that the non-prioritized car did not stop safely before encroaching upon the prioritized lane. Therefore, the first condition for safe distance in this situation does not hold.

The next step would be to examine whether or not one of the other two conditions hold, in light of the scenario right before the dangerous situation was triggered. If before the dangerous situation only the longitudinal distance was safe, it indicates that the non-prioritized car encroached upon the prioritized lane at a point where the oncoming car had sufficient time to prevent a collision, as depicted in Figure 17(a). In this case, the prioritized car is to blame for not following the above braking pattern as stated in the proper response to this situation. In a case where only the lateral distance was safe right before the dangerous situation, it implies that the oncoming car did not have sufficient distance to prevent a collision according to condition 2. The proper response to this situation is for both cars to brake laterally. Since the oncoming car has no lateral velocity as it is going straight, only the non-prioritized car has to brake laterally. Hence, in the case of a collision the non-prioritized car would be to blame.

Therefore, we should examine whether the third condition holds according to the state right before the dangerous situation is triggered, using the same methodology demonstrated in the previous cases.




Figure 17(a): this is a situation where only the lateral distance is non-safe. The dangerous situation will be initiated when the longitudinal distance also becomes non-safe





Figure 17(b): a dangerous situation due to non-safe lateral and longitudinal distance. If prior to this situation only the lateral distance was safe it means that the red car maneuvered laterally into the blue car's path, leaving it without sufficient time to prevent a collision

Vulnerable Road Users (VRUs) and Occlusion Scenarios

Scenarios as depicted in the NHTSA research:



VRU crash without prior vehicle maneuver

In the cases depicted above, a vehicle encounters a VRU at a straight non-junction road. In order to determine who is responsible for such a collision we first need to examine whether the vehicle was driving at an unreasonable speed or not, with respect to the speed limit, weather conditions and obstructions. Assuming the driving speed was reasonable, we examine whether the two conditions we previously defined hold.

Consider a situation as depicted in Figure 18. The car should not be held responsible in case it did not accelerate during the response time (the time interval that starts at the exposure time until the end of the response time) and performed a longitudinal brake until the collision or until achieving a full stop. On top of that, the average speed of the car in the time interval between the exposure time until the collision has to be smaller than the average speed of the pedestrian in the same time interval. In a situation where both of these conditions hold, we can see that even if the driver was driving at a velocity below v_{limit} , he still did not have a chance to prevent the collision. When the pedestrian was first revealed to the driver it was already in a non-safe longitudinal distance and, therefore, avoiding a collision in such a time span is not feasible. In other words, we can say that the car maintained sufficiently low speed, but the pedestrian's average speed, with respect to his lateral distance from the car, is above a reasonable threshold, thus the pedestrian is at fault.

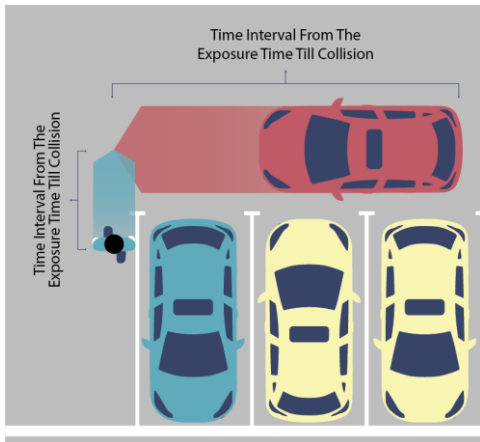


Figure 18: a pedestrian is revealed to the car at the exposure time. After the response time the car brakes until colliding with the pedestrian or reaching a full stop

Conclusion

In the above analysis, we have illustrated how the RSS model covers all the multi-agent scenarios presented in the NHTSA research (excluding only vehicle failure events, crashes with stationary objects etc.). Regarding scenarios related to unstructured roads, the RSS model does provide comprehensive coverage but there is not enough time and space in this article to address these areas. We aim to address them in a further article.

The NHTSA typology provided us with an efficient platform for demonstrating the soundness of the RSS model. **Indeed, our evaluation of the model with respect to real-world scenarios only reinforces the need to create this type of comprehensive, useful, and transparent policy that allows autonomous vehicles to cope with the unpredictable nature of human drivers. While at the same time providing a mechanism to clarify liability in the case of collisions between AVs and human drivers.** The combination of reducing risk of machine error and clarifying liability in a conclusive manner is critical to allow the industry to move forward in the development of AVs. AV developers should be held liable for errors of the machine's responsibility, but these must be held to an absolute minimum in order to build consumer trust. At the same time, we must not allow innovation aimed at bringing massive benefits to society to be stifled by liability risk in cases where the machine acted properly under the rules of human judgment. **We believe the industry can come together to create a collaborative platform which provides a "safety seal" that first and foremost will create a safer product, but at the same time will protect OEMs from unreasonable and unwarranted liability assignment by regulators and society.** For this to succeed, all industry players must join forces - from OEMs to suppliers to standards-bodies to regulatory bodies - especially to ensure that this model does not benefit a certain technology or vendor. We are putting a stake in the ground in an attempt to drive the industry to agree there is a definitive need to formalize the rules of judgment, responsibility and fault in order to prepare the landscape for successful deployment of AVs, thus realizing the massive benefits to society. Far from a system designed to avoid liability, RSS is an innovative model intended to enable AVs to perform to the highest safety standards. **The Responsibility-Sensitive Safety model is a viable place to start the conversation of how to achieve a shared greater goal: profound safety in the autonomous ecosystem.**