WISE Requirements Analysis Framework for
Automated Driving Systems

Operational Design Domain for Automated Driving Systems

Taxonomy of Basic Terms

Krzysztof Czarnecki
Waterloo Intelligent Systems Engineering (WISE) Lab
University of Waterloo
Canada

July 21, 2018

For updates and related documents, see
https://uwaterloo.ca/waterloo-intelligent-systems-engineering-lab/projects/wise-requirements-analysis-framework-automated-driving

Document history

Updates	*Date*	*Contributor(s)*
Version 1 created	October 23, 2018	Krzysztof Czarnecki[1]
Version 1.1 Editorial changes	July 21, 2018	Krzysztof Czarnecki[1]

[1]Waterloo Intelligent Systems Engineering Lab
University of Waterloo

Abstract

This document defines a taxonomy of basic terms used in the description of an Operational Design Domain (ODD) for an Automated Driving System (ADS). Among others, the taxonomy defines operational world models and terms for specifying driving scenarios and their attributes.

Table of Contents

1. Scope

This document defines a taxonomy of basic terms used in the description of an Operational Design Domain (ODD) for an Automated Driving System (ADS). These terms provide a basis for creating operational world models, which are used to specify driving tasks and requirements and to verify and validate the implemented ADS behavior.

The taxonomy targets SAE J3016 levels of driving automation 4 and 5. It can also be used for lower levels of driving automation; however, the current scope of this document does not consider interaction with a driver or a fallback-ready user.

The document was developed with the intent to be consistent with prior art to the extent possible, including relevant industry standards and scientific literature. The following format is used when referring to terms defined in the industry guidance documents: [XXX](N.M….), where [XXX] references a document in Section "References", and (N.M….) is the definition number defined in the cited document. For example, [FuS](1.103) refers to definition 1.103 in ISO 26262:2011.

2. Basic Terms

2.1 Operational Design Domain (ODD)

J3016 defines an Operational Design Domain (ODD) as "operating conditions under which a given driving automation system or feature thereof is specifically designed to function, including, but not limited to, environmental, geographical, and time-of-day restrictions, and/or the requisite presence or absence of certain traffic or roadway characteristics." [LA](3.22) ODD is the design domain of an ADS or a feature thereof with respect to its operation. J3016 introduced this concept in order to capture limitations for driving automation at levels 1, 2, 3 and 4. Level 5 ADS (full driving automation) has an unlimited ODD, which offers the same mobility as a human driver.

An ODD may put limitations on

1. the road environment
2. the behavior of the ADS-equipped subject vehicle; and
3. state of the vehicle.

Generic road-environment limitations may include types of roads (urban, rural, or freeways), specific elements such as roundabouts or tunnels, temporary structures such as construction zones, traffic volumes, and weather and visibility conditions. Road-environment limitations may also be specific, such as a geographical area, or season, or time of the day.

Road-environment limitations are the most common elements of an ODD, but an ODD may also place limitations on ADS behavior and state. Limitations on ADS behavior may include speed limitations or limitations on maneuvers, such as reversing. Limitations on the state of the subject vehicle may include the requirement of no trailer being attached, loading limitations, or minimum tire inflation level.

The ODD may reflect the requirements of a particular driving automation feature. J3016 defines a driving automation system feature as "a level 1-5 driving automation system's design-specific functionality at a given level of driving automation within a particular ODD, if applicable" [LA](3.9) Examples of driving modes include low-speed traffic jam assist, high-speed cruise assist, valet parking, and urban in-lane driving. J3016 stipulates that an ADS may have one or more features, and each feature has exactly one ODD.

2.2 Operational Road Environment Model

The design and verification and validation of an ADS requires modeling the operational environments in which the ADS is to operate the ADS-equipped vehicle.

An operational road environment model (OREM) is a representation of the relevant assumptions about the road environment in which an ADS will operate the ADS-equipped vehicle.

Environment models capture the properties of the environment that are relevant to the ADS operation, while abstracting irrelevant details. OREMs can represent generic environments, such as two-lane rural road, or actual roads in a specific geographic area. OREMs can take different forms, including specification documents and executable models. While a closed course used for testing may be designed to model a target deployment environment, within this document, OREMs are meant to be documents or software models. OREMs provide context for specifying driving tasks of an ADS and for verifying and validating the ADS. Executable OREMs can be used in simulation testing. OREMs can also specify the environment for closed course and field tests.

An ODD of an ADS implies a set of operational environments in which the ADS can operate the ADS-equipped vehicle. These environments can be specified using a set of OREMs. OREMs may be *fully* or *partially in scope* of an ODD, or they may be *out of scope* of an ODD. For example, an OREM for a highway pilot may include a section of a highway, ramps, and parts of roads accessing the ramps. The highway and the ramps may be in scope of the highway pilot's ODD, but the access roads may be outside of the ODD. Including the parts of the roads accessing the ramps in the OREM allows specifying the ODD boundary. Further, out-of-scope OREMs are also useful to specify ODD boundaries and may be used to test the capability of the ADS to detect an out-of-ODD condition and the fallback performance in the event of ODD departure.

2.3 Subject Vehicle Model

Subject vehicle model represents the vehicle operated by the ADS under development. The model includes the aspects of the subject vehicle systems that are relevant to the ADS at a level of fidelity required for the analysis at hand. These are typically vehicle dynamics, including steering, suspension, and braking systems, and powertrain (J3016 refers to them as "other vehicle systems" [LA]).

2.4 Operational World Model

An operational world model (OWM) is a composition of an OREM with one or more subject vehicle models, each coupled with a model of a subject ADS or an actual subject ADS, that is, an ADS under development.

An operational world model (OWM) supports requirements specification and verification and validation of an ADS. An OWM must be distinguished from the reflective *ADS world model*, which is an representation of the world within an ADS. The J3131 draft defines the *ADS world* *model* as “the ADS’s internal representation of the portions of the environment that the system is aware of or that are of interest to the system and the user for the purpose of achieving a specific goal. The world model includes models for the system itself and entities in the environment.” Thus, the ADS world model is inherently partial with respect to the operational world. Also, in contrast to OWM, elements of the ADS world model are typically associated with levels of confidence about their validity with respect to the real world (or OWM in simulation).

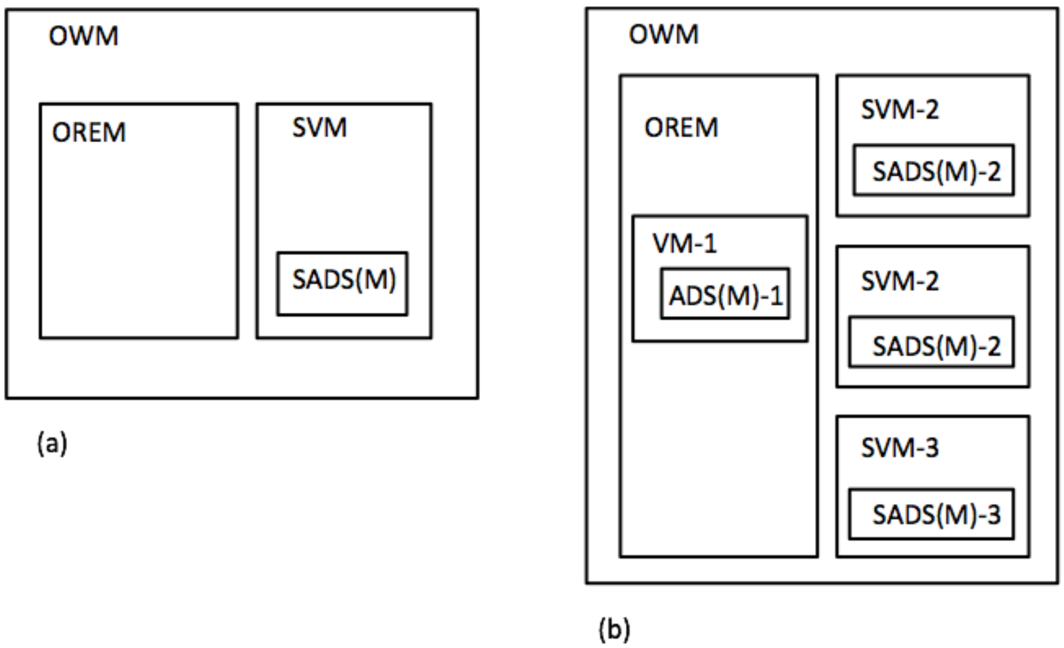

Figure 1 Components of an Operational world model (OWM); (a) OWM with a single subject vehicle model (SVM) and subject ADS or subject ADS model; (b) OWM with multiple subject vehicle models (SVM-2,-3) and subject ADSs or subject ADS models (SADS(M)-2,-3); note that the OREM may include one or more vehicle models that are not subject to development

An OWM may include a single subject vehicle (Figure 1a) or multiple subject vehicles equipped with the same subject ADS under development, or multiple different subject ADSs under development (Figure 1b). The latter may be necessary to, for example, design the interaction between ADS-operated low-speed buses and ADS-operated passenger vehicles, where both types of ADS are in scope of the development. The OREM of the OWM may also include other human-operated vehicles and vehicles equipped with other existing ADSs that are not under development (ADS(M)-1 in Figure 1b).

The representation of a subject ADS that is included in an OWM depends on the type of analysis being performed. The OWM contains (1) a *model* of the subject ADS in the case of model-in-the-loop testing, (2) the subject ADS *software* in the case of software-in-the-loop testing, and (3) the subject ADS *software and hardware* in the case of hardware-in-the loop testing.

An OWM defines *world state* that evolves over time. The state includes entities and their attributes and relationships and any state derived through additional rules or functions.

An ODD may be specified using multiple OWMs, which characterize the environments that the subject ADS can operate the subject vehicle in and any restrictions on the ADS. For example, separate OWMs could be specified for urban driving and highway driving. Or each geographic area may have its own OWM defined.

2.5 Scene and Situation

> *A scene is a snapshot of the world state, including road configuration, static objects, dynamic objects, and environmental conditions, but excluding unobservable state.*
>
> *A situation is a scene extended with behaviors and observable and unobservable states of all dynamic objects, i.e., road vehicles (including operators), pedestrians, and animals. The state of a dynamic object includes relevant unobservable state, such as intent (e.g., destination and other agent's mission objectives), the actions available to the agents, their perception of the world, and the values driving their decisions (e.g., traffic rules, social norms).*

These definitions of scene and situation are inspired by existing literature [GBFR14,UMR15]; in particular, Ulbrich et al. give a consolidated and in-depth analysis of these terms.

We assume that the OWM includes models behavioral models of each agent, and thus the snapshot also represents each agent's subjective representation of the world. Thus, snapshots of the OWM are situations. An important observation is that the OWM can only be complete and accurate in simulation.

An OWM will normally be supplemented with a catalog of explicitly specified situations, which will be used to specify and verify an ADS. The situations play the role of examples in the spirit of example-driven modeling [BZC14].

A situation specification may assume an open or closed world [BDA13]. Under the *open-world assumption*, elements may be present in the modeled world irrespectively of whether they are actually included in the specification or not. For example, an open-world situation specification may represent an intersection with two vehicles; this means that the modeled world has at least the two vehicles, but the absence of additional vehicles on the modeled intersection cannot be concluded based on this specification. Under the *closed-world assumption*, all elements that are present in the modeled world are also included in the specification and vice versa. In our example, we would conclude that just the two vehicles and no other vehicles are present at the intersection.

Situation specifications can also be partial or complete [BDA13]. A *complete specification* fixes the complete world state; a *partial specification* fixes only part of the state. Open-world models are inherently partial: they need to be closed before the scope of the complete state is known. In our example, closing the open-world specification of the intersection would mean committing to a particular number of other vehicles, such as zero or more. Closed-world specifications can be partial or complete [BDA13]. A partial closed-world specification may have parameterized presence of elements and also parameterized element properties. For example, a partial closed-world model of an intersection may stipulate the presence of between one and four vehicles, including pose and speed ranges of each vehicle. Completing this specification would require committing to a particular number of vehicles between one and four and selecting a pose and speed of each vehicle within the valid range. Partial specifications of situations can also be *probabilistic*, where the open parameters are constrained by probability distributions or probability density functions. In our example, the number of vehicles between zero and four could be constrained by a probability distributions and the vehicle poses and speeds would be constrained by probability density functions.

Situations may also involve *temporal abstraction*. When specifying behavior at a higher level of abstraction, such as tactical as opposed to operational, multiple snapshots may be collapsed into a single situation. For example, a particular behavior specification may treat the approach to an intersection as a single situation, even though, at a more detailed level, there are multiple snapshots with the vehicle getting closer to the intersection in each snapshot. In other words, the state evolution within the duration of a situation is abstracted. An abstraction needs to provide necessary and sufficient conditions for a situation to hold, which allow to reason about transitions among situations.

2.6 Scenario and Story

> *A scenario represents an evolution of the world state of an operational world model over time.*

In a discrete time model, a scenario is simply a time-stamped sequence of scenes or situations.

Whenever relevant, a distinction between a *scenario* as sequence of scenes and a *story* as a sequence of situation can be made. Unless required for clarity, we use the term scene to cover both meanings.

Similar to a situation specification, a scenario specification can also have open or closed world interpretation and be partial or complete. A complete scenario specification would provide the complete state for every time instant. A partial scenario specification may specify some of the state partially. Partial scenarios can be represented by probabilistic behavioral models.

2.7 Situation and Scenario Occurrence and Property Statistics

Situations and scenarios should be supplemented with statistics characterizing chances of their occurrence and the distributions of their properties in the real-world operational road environment. The chances of occurrence are measured using *rates* or *proportions* with exposure in the denominator.

Exposure of a situation or scenario is an event that provides opportunity for the situation or scenario to occur [Elv08,EEC09]. *Units of exposure* are either discrete or continuous. An example of a discrete unit of exposure for the situation of violating a red light is entry onto a signalized intersection. Examples of continuous units of exposure are *vehicle distance travelled* (e.g., in kilometers) or *vehicle time traveled* (e.g., in hours).

The chances of occurrence of a situation or scenario can then be expressed as *frequencies*, if the situation or scenario occurrence count is relevant; or *proportions* of traveled distance or time, if the cumulative situation or scenario duration is relevant.

Relative frequencies are rates between situation or scenario occurrence count and exposure event count. Examples include number of red light violations per signalized intersection entries, and number of signalized intersection entries per mission, such as a trip. Relative frequencies can also be interpreted as *probabilities*, where the unit of exposure would represent independent trials; however, one would still need to specify the unit of exposure when reporting such probabilities.

Temporal and spatial frequencies use duration or distance traveled in their denominator. *Temporal frequencies* count situation or scenario occurrence per vehicle time traveled, such as the number of intersection entries per hour traveled. *Spatial frequencies* count situation or scenario occurrence per vehicle distance traveled, such as the number of intersection entries per kilometer traveled.

Proportion of time or distance traveled specifies the proportion of the time or distance traveled in a given situation or scenario. These two metrics are used for situations and scenarios that span significant portions of a trip. Examples include proportion of distance travelled in the left-most lane or proportion of time travelled at night. For some situations and scenarios, such as traversing an intersection, it may be useful to provide both frequencies and proportions of time or distance travelled.

Situations and scenarios may represent exposure for other situations and scenarios. In particular, situations that provide opportunities for crashes are *crash exposures*. For example, red light violation is exposure for a collision. When performing hazard analysis and risk assessment (HARA), situations that provide opportunities for hazardous events are analyzed and the chances of occurrence of these situations are assessed. ISO 26262 provides an interval-based scheme to classify the occurrence measures of exposure situations consisting of four classes, E1, E2, E3, and E4. E1 represents rare events, and E4 represents events that occur most of the time when operating a vehicle. These classes of occurrence of exposure can be frequency or duration based, and the E1-4 interval scheme simplifies the estimation and analysis by reducing the burden of having to provide precise statistics. Appendix B of ISO 26262, Part 3, gives guidelines and examples for using the classification scheme. In particular, the guidance uses two metrics: (1) percentage of operating time and (2) frequency of occurrence per year. Estimating the occurrence rate of rare events such as crash situations is challenging and may benefit from special statistical tools such as extreme value theory (see Section 2.7.3).

Selecting adequate units of exposure for crash situations and scenarios is challenging. The most widely used unit of exposure for vehicle crashes is vehicle distance travelled, which leads to the rate of crashes per vehicle distance travelled. The statistic does not take into account relevant aspects such as the driver learning effect. Consequently, Elvik et al. propose elementary units of exposure that can be used for most crash situations [Elv08,EEC09]. These units are discrete event types including opposite traffic encounters, simultaneous arrivals at conflict points, changes of travel direction close to other road users, and braking or stopping. The only continuous exposure is distance travelled for road departures, resulting in road departure frequency per kilometers driven. Elvik et al. demonstrate how to estimate the occurrence rates of the discrete units of exposure from *annual average daily traffic* (AADT).

Partial situations may be composed in parallel, and the corresponding occurrence rates may be multiplied when the situations are independent. For example, if we assume that night driving and traversing intersections are independent, that is, the chances of crossing an intersections are the same whether the driving is done at night or day, then the proportion of night driving can be multiplied with the rate of intersection traversals in order to obtain the rate of intersections traversed at night with respect to total driving.

Statistics of various other properties of situations and scenarios are also relevant to modeling the operational environment. These include duration of situations and scenarios, gap acceptance, numbers and speeds of different road users, etc. The probability distributions are also exposure based, e.g., with units of exposure being independent trials. An example would be the distribution of the number of vehicles at an intersection per intersection entry. The probability distributions of these random variables need to be established as part of the environment modeling effort.

2.8 Risk Category and Risk Measures

Each situation and scenario is classified according to its *risk category*. The risk category is *normal driving*, *near crash*, *crash*, or *fallback*. *Crash situations and scenarios* are those in which a crash cannot be avoided. *Near-crash situations and scenarios* are those where a crash is imminent but still can be avoided. For example, near-collision situations and scenarios require an emergency maneuver of one or more of the involved road users in order to avoid the collision. All other situations and scenarios are *normal driving ones*. Hyden summarized the three categories of situations and scenarios as a pyramid, reflecting the fact that crashes and especially severe crashes are rare (Figure 2). For example, 0.5% of all police-reported crashes were fatal in the US in 2015 [TSF16]. *Fallback* situations and scenarios are those in which the ADS is about to leave its ODD or experiences a failure that necessitates a fallback maneuver [LA].

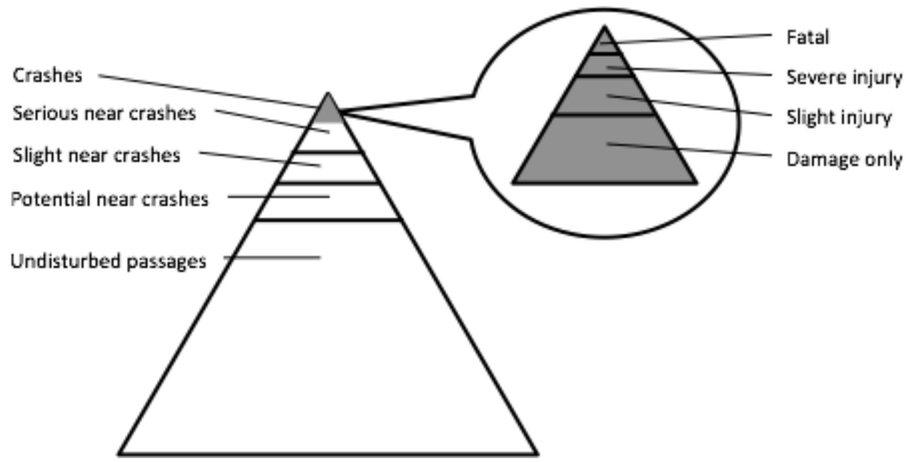

Figure 2 Pyramid concept (adapted from [Hyd87])

Within each risk category, a scene can be assessed for its risk level. The risk level of a scene is based on *exposure* and *severity* measures: the scene is exposure for a range of possible crashes of different severity. For example, entering an intersection on red light to make a through movement has associated risk of a front-to-side or angle collision, with severity related to the operating speed of vehicles at that intersection. The risk level assessment would need to consider the distribution of subsequent situations until the intersection traversal is complete or results in a crash. A third factor that enters the risk level assessment is *controllability*, which is defined as the “ability to avoid a specified harm or damage through the timely reactions of the persons involved, possibly with support from external measures” [FuS]. The key person considered for controllability in the case of a human-operated vehicle is the human driver. For an ADS-operated vehicle, there are still other persons to be considered for controllability assessment, such as drivers of human-driven vehicles interacting with the subject vehicle, fallback-ready users (who may be remote), and pedestrians and cyclists. For example, a human driver of a vehicle that follows an ADS-operated vehicle may be able to react timely to a sudden braking of the ADS-operated vehicle to avoid a rear-end collision.

Similarly, scenarios are classified according to risk level, where the scenario is considered exposure for a range of possible crashes of different severity. A scenario may consist of entering an intersection on red light to make a through movement, followed by an evasive maneuver to avoid a collision. The risk level of the overall scenario considers the specific sequence of the two situations, which is a more specific behavior than traversing an intersection on red light (which may or may not involve a conflict with another vehicle).

2.9.1 Crash Severity

Crash situations and scenarios have associated *crash severity*, which assesses the level of losses resulting from the crash. In general, losses include personal injury or death or property damage. Severity assessment is a complex matter in practice. In addition to direct injuries and deaths occurring in a crash, consequences of a crash to human health and life may reach far beyond the crash occurrence. For example, a crash may cause an undiagnosed aneurysm in the brain of a person involved in the crash, which may lead to a sudden brain hemorrhage and the person's disability or death years after the crash. Similarly, in addition to *direct property losses*, such as vehicle and infrastructure damage, and *indirect economic losses*, such as impact on employment and productivity, there are also *non-economic losses*, such as reduced quality of life for the injured and his or her relatives.

ISO 26262 focuses on losses that represent personal injury, which includes fatal injuries (deaths) and non-fatal ones, and it does not consider property damage. The standard provides guidance on the use of different injury severity scores, such as the five-level Abbreviated Injury Scale (AIS). The AIS is maintained by the Association for the Advancement of Automotive Medicine (AAAM), which publishes a classification guide with examples for each of the seven AIS levels [AIS15]:

- AIS 0: no injuries;
- AIS 1: light injuries such as skin-deep wounds, muscle pains, whiplash, etc.;
- AIS 2: moderate injuries such as deep flesh wounds, concussion with up to 15 minutes of unconsciousness, uncomplicated long bone fractures, etc.;
- AIS 3: severe but not life-threatening injuries such as skull fractures without brain injury, spinal dislocations below the fourth cervical vertebra without damage to the spinal cord, etc.;
- AIS 4: severe injuries (life-threatening, survival probable) such as concussion with up to 12 hours of unconsciousness, paradoxical breathing;
- AIS 5: critical injuries (life-threatening, survival uncertain) such as spinal fractures below the fourth cervical vertebra with damage to the spinal cord, intestinal tears, cardiac tears;
- AIS 6: extremely critical or fatal injuries such as fractures of the cervical vertebrae above the third cervical vertebra with damage to the spinal cord, extremely critical open wounds of body cavities.

AIS classifies single injuries to a person; in the case of multiple injuries to a person, Maximum AIS (MAIS) or Injury Severity Index (ISI) [BOH74, BOM00] may be used.

Crash statistics allow determining the *distribution of injuries* that can be expected to occur in different types of crashes (e.g., [Kwe11, Abd17, YPA08, GBA14]). The key idea is to derive a statistical fit between physical parameters of crashes and the injury distribution. For example, the *Delta-V* metric, which is the change of a velocity vector experienced by a road user during a crash [LDK17], allows estimating the occupant injury severity in frontal crashes [GG08, KG12]. Other relevant crash parameters affecting collision severity are impact point and impact angle. For road departures and rollovers, the key parameters are vehicle speed and terrain structure at the crash site. Rollovers are the most injurious types of accidents compared to any other accidents [CK05]. An analysis of historical rollover data identified the following significant factors that correlate with occupant injury level: the number of roof impacts, impact with fixed objects prior rollover, rollovers stopped by an impact with a fixed object, and failing to wear seatbelts [DE04].

ISO 26262 provides a *four-level severity classification* of hazards, and specifically, *crash types* (see Table B.1 in [FuS]). The classification sets bounds on AIS levels and probability of occurrence of injuries at these levels. S0 corresponds to A0 and less than 10% probability of AIS 1-6. S1 corresponds to more than 10% probability of AIS 1-6 (and not S2 or S3). S2 corresponds to more than 10% probability of AIS 3-6 (and not S3). S3 corresponds to more than 10% probability of AIS 5-6. Table B.1 in [FuS] provides examples of crash types for each severity class.

A crash may result in multiple individuals injured and a range of property losses. It may be possible to meaningfully aggregate different losses by their kind and severity. For example, aggregate loss from a crash could be characterized by stating the number of individuals suffering injuries at each of the seven MAIS levels and the total value of property damage. In that case, the risk of a given crash type would be the joint probability distribution over these eight random variables. For some applications, such as comparing the severity of different crash types across road-network locations using summary statistics, all losses can be combined using the *loss value method*, which computes a total loss as a weighted sum of individual losses [CK05]. In essence, the method assigns specific monetary cost to each injury type and property damage.

2.9.2 Near-Crash Situations and Scenarios

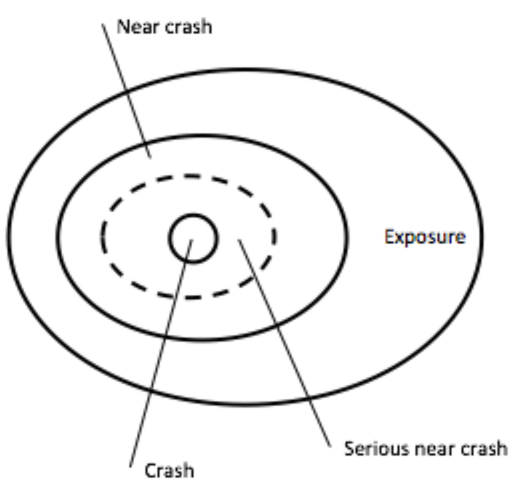

Figure 3 Crash, near-crash, exposure (adapted from [AH77])

Near-crash situations and scenarios are those where a crash is imminent but still can be avoided. They require an emergency maneuver by one or more involved road users in order to avoid a crash. The maneuver may be emergency braking or steering for a vehicle, or a sudden evasive movement of a pedestrian. Other near-crash situations and scenarios are instances of skidding where avoiding a crash requires skid recovery, and near rollovers that require specific ADS inputs to avoid a crash, such as reducing lateral force or avoiding a trigger. Near-crash situations and scenarios are exposure events for crashes that are closest to the crashes (Figure 3), where the proximity is defined using surrogate safety measures and specific thresholds.

2.9.3 Surrogate Safety Measures

Near-crash situations and scenarios may be identified using *surrogate safety measures*, also known as *proximal safety measures* [MFH17]. In general, such measures can be classified as describing proximity in time, proximity in distance, intensity of the necessary evasive action, and other events, which may be situation specific, such as red-phase violations, stop-bar encroachments, high speed differentials when merging, and distribution of merge points.

The most commonly used safety surrogate measure is *time-to-collision (TTC)*. TTC represents the urgency to brake in order to avoid a collision [Lee76]. TTC can be computed with the assumption of constant speed or constant acceleration and is used mainly to identify situations that may lead to rear-end collisions and collisions with pedestrian and stationary objects. J2944 provides guidance on computing TTC depending on application context and cites studies on TTC distributions found in naturalistic studies and thresholds indicating high-risk situations [DP15]; Schwartz provides a very general and efficient method to compute TTC [Sch14].

Other well-studied surrogate safety measures include *post-encroachment time (PET)*, *time gap*, and *required deceleration rate* (RDR) [DP15, GH03]. Time gap is the time lapse until a vehicle arrives at a conflict point assuming constant speed; PET is the time laps until a vehicle actually arrives at a conflict point. PET and time gap are used for many types of situations, including rear-end and at angle approaches such as car following and merging and crossing traffic at intersections. RDR makes the level of braking needed to avoid a collision explicit and is used in similar situations as TTC and PET. J2944 provides guidance on computing PET and RDR, their distributions in naturalistic driving, and safety-relevant thresholds.

NHTSA published two reports summarizing results of their studies evaluating surrogate safety measures using simulation and field data [GH03, GPS08]. The measures are complementary and may need to be used together, potentially composed into surrogate safety indices. For example, TTC is infinite when following a front vehicle at equal speed, even though the following time gap may be too small to avoid a collision if the front vehicle was to suddenly brake. Thus, a minimum time gap needs to be observed. Further, time gap is infinite at standstill, yet a minimum distance gap at standstill should be observed. Mahmud et al. survey a total of 38 surrogate measures and indices, summarizing their applications, and advantages and limitations [MFH17].

For near-crash *scenarios*, summary measures are used. *Minimum TTC* is the minimum value of TTC reached at any point during the scenario. *Time Integrated TTC (TITTC)* is a weighted time interval over which the TTC is less than some exposure threshold, with times weighted by how far below that threshold the TTC is at each moment. J2944 provides guidance on computing both summary measures [DP15].

Since crashes are rare events, crash data is scarce by definition. Near-crashes are used as surrogate events for crashes to estimate crash rates. A popular practice is to assume a *constant conversion rate* from near crashes to crashes [WJ12]. For example, the NHTSA study of surrogate measures found a conversion rate of 20,000 near crashes to one crash [GPS08]. However, in general, the conversion rate depends on situation [GKM10] and the relation between crashes and near crashes is non-linear. A more accurate method is to use *extreme value statistics* [ST06]. This methods relies on the idea that the distribution of extreme values of a random variable are independent of the underlying distribution of the variable and follow a generalized extreme value distribution, such as the generalized Pareto distribution. By fitting the generalized extreme value distribution to the observed extreme values of a safety surrogate measure, the rate of the unobserved extremes of larger magnitude, which correspond to crashes, can be estimated. Consequently, the method allows estimating a more accurate crash rate than Poisson statistics using the same amount of field data. Extreme value theory has been applied to safety analysis related to intersection collisions [ST06], lane departures [Tar12], lane change collisions [ZIM14], and passing maneuver collisions [FA17].

2.10 Situation and Scenario Demand

Situations may place a different *demand* on a road user, including the ADS, depending on situation complexity and environmental factors. Situations are classified into *low*, *medium*, *high*, *very high*, and *extreme demand*. Driving in good weather and lighting, in low traffic volumes, and at low speed corresponds to low demand. Demand is increased by, among others, higher speeds, lower gap acceptance, poor weather, poor visibility, occlusions, slippery roads, high traffic volumes, temporary structure such as temporary lane assignments, and complex urban environment with unpredictable pedestrians and cyclists. Higher demand normally requires higher on-road performance of the road users. Note that demand and risk category are separate dimensions. Normal driving situations can have varying levels of demand, and so do near-crash and crash scenarios.

Scenarios can also be classified according to demand, but different situations in a scenario may have different demand.

According to *task difficulty homeostasis* theory, driving task difficulty results from the interaction between situation demand and the road user performance or capability [Ful05]. Fuller argues that human drivers target a specific level of task difficulty that they are comfortable with and that this choice determines their driving behavior, such as their speed choice and gap acceptance. He provides evidence that task difficulty and perceived risk appear to be very highly correlated, but perceived risk and ratings of statistical risk are completely unrelated until a critical speed is reached (presumably where task demand approaches driver capability). The statistical risk of collisions increases sharply when the situation demand surpasses the road user capability.

While the demand concept was originally proposed for explaining human driver behavior, it is still applicable to an ADS. In particular, higher demands may require higher computational resources, for example, to allow processing a higher number of relevant elements in the environments within short periods of time or the ability to deal with increased perceptual noise. Demanding situations can be used to stress test an ADS. Demands that exceed system capabilities imply ODD departure. However, the situational elements that increase the demand for a human driver vs. an ADS are likely different. Thus, the demand of different situations needs to be assessed with respect to the functional concept of an ADS, including its planned sensing capabilities. Just as a human driver, an ADS may regulate demand by choosing adequate speed and avoiding situations that may exceed its capability.

2.11 Restricted Operational Domain

> *The Restricted Operational Domain (ROD) is the specific conditions under which a given driving automation system or feature thereof is currently able to function* [CP18].

The ROD is a system degradation concept, which allows restricting the ODD of a driving automation system or its feature based on the current capability of the system. For example, an ODD of an ADS may include both urban roads and freeways; however, failure of a long-range radar may restrict its capability to driving only on roads with maximum speed of 50 km/h. As a result, the ROD of the system would be urban roads with speed limits of 50 km/h or lower.

2.12 Operational World Model Ontology

> *Operational world model ontology is a conceptualization of the elements that occur in an operational world model. The conceptualization includes element types, element attributes and relationships and, if applicable, behaviors.*

An OWM ontology contains concepts in five categories:

1. *Road structure*: This category includes road geometry, lane configuration, roadside structure, traffic control devices, junctions, and temporary structure.
2. *Road users*: This category includes vehicles and pedestrians and their behavior.
3. *Animals*: This category covers wild and domestic animals of different sizes.
4. *Other obstacles*: This category covers all other obstacles that might be found on a roadway, such as lost cargo, tree branches, or debris.
5. *Environmental conditions*: This category consists of atmospheric, lighting, and road surface conditions.

The ontology is defined elsewhere [O18a,O18b].

References

[Abd17] A. Abdulhafedh. Road crash prediction models: different statistical modeling approaches. Journal of transportation technologies. 2017 Mar 13;7(02):190

[AH77] F.H. Amundsen and C. Hyden, Eds. Proceeding of First Workshop on Traffic Conflicts. Institute of Transport Economics. Oslo/Lund Institute of Technology, Oslo, Norway, 1977

[AIS15] Abbreviated injury scale; Association of the advancement of Automotive medicine; Barrington, IL, USA, https://www.aaam.org/abbreviated-injury-scale-ais/

[BDA13] K. Bąk, Z. Diskin, M. Antkiewicz, K. Czarnecki, and A. Wąsowski, Partial Instances via Subclassing. In Proceedings of the 6th International Conference on Software Language Engineering (SLE), Indianapolis, USA, 2013

[BOH74] S.P. Barker, B. O'Neill, W. Haddon, W.B. Long. The injury severity score: a method for describing patients with multiple injuries and evaluating emergency care. The Journal of Trauma, Vol. 14, No. 3, 1974

[BOM00] Z. Balogh, P.J. Offner, E.E. Moore, W.L. Biffl. NISS predicts post injury multiple organ failure better than ISS. The Journal of Trauma, Vol. 48, No. 4, 2000

[BZC14] K. Bak, D. Zayan, K. Czarnecki, M. Antkiewicz, Z. Diskin, A. Wasowski, and D. Rayside. Example-Driven Modeling: Model = Abstractions + Examples. In ICSE 2014, New Ideas and Emerging Results (NIER)

[CK05] J.R. Campbell and K.K. Knapp. Alternative crash severity ranking measures and the implications on crash severity ranking procedures. In Proceedings of the 2005 Mid-Continent Transportation Research Symposium, 2005

[CP18] I. Colwell, B. Phan, S. Saleem, R. Salay, and K. Czarnecki. An Automated Vehicle Safety Concept Based on Runtime Restriction of the Operational Design Domain. In Proceedings of Intelligent Vehicles Symposium, 2018

[DE04] K.H. Digges and A.M. Eigen. Crash attributes that influence the severity of rollover crashes. Progress in Technology, 2004;101:299-308

[DP15] Operational Definitions of Driving Performance Measures and Statistics. SAE Surface Vehicle Recommended Practice, J2944:2015-06

[EEC09] R. Elvik, A. Erke, P. Christensen. Elementary units of exposure. Transportation Research Record: Journal of the Transportation Research Board. 2009 Sep 4(2103):25-31

[Elv08] R. Elvik. Towards a general theory of the relationship between exposure and risk. Transportøkonomisk Institutt, TØI report 1316/2014, Olso, Norway, 2014

[FA17] H. Farah and C.L. Azevedo. Safety analysis of passing maneuvers using extreme value theory. IATSS Research. 2017 Apr 30;41(1):12-21

[Ful05] R. Fuller. Towards a general theory of driver behaviour. Accident Analysis & Prevention. 2005 May 31;37(3):461-72

[FuS] Road vehicles — Functional safety. ISO 26262:2011

[GBA14] R. Garrido, A. Bastos, A. de Almeida, and J.P. Elvas. Prediction of road accident severity using the ordered probit model. Transportation Research Procedia. 2014 Jan 1;3:214-23.

[GG08] D.J. Gabauer and H.C. Gabler. Comparison of roadside crash injury metrics using event data recorders. Accident Analysis & Prevention. 2008 Mar 31;40(2):548-58

[GH03] D. Gettman and L. Head. Surrogate safety measures from traffic simulation models: final report. Federal Highway Administration, FHWA-RD-03-050, 2003

[GKM10] F. Guo, S.G. Klauer, M.T. McGill, and T.A. Dingus. Evaluating the relationship between near-crashes and crashes: Can near-crashes serve as a surrogate safety metric for crashes?. DOT HS 811 382, Virginia Tech Transportation Institute, 2010

[GPS08] D. Gettman, L. Pu, T. Sayed, and S.G. Shelby. Surrogate safety assessment model and validation: final report. Federal Highway Administration, FHWA-HRt-08-051, 2008 Jun

[Hyd87] C. Hyden. The Decelopment of Method for TrafJic Safety Ecaluntion: The Swedish Traffic Conflict Technique. Bull. 70. Lund Institute of Technology, Lund, 1987

[KG12] K. D. Kusano and H. C. Gabler. Safety Benefits of Forward Collision Warning, Brake Assist, and Autonomous Braking Systems in Rear-End Collisions. IEEE Trans. Intell. Transp. Syst., vol. 13, no. 4, pp. 1546– 1555, Dec. 2012

[Kwe11] Y.-J. Kweon. Crash Data Sets and Analysis. Chapter 8 in B.E. Porter (Ed.) Handbook of traffic psychology. Academic Press, 2011

[LA] Surface Vehicle Recommended Practice — Taxonomy and Definitions for Terms Related to Driving Automation Systems for On-Road Motor Vehicles. SAE J3016:JUL2018

[LDK17] A. Laureshyn, T. De Ceunynck, C. Karlsson, Å. Svensson, and S. Daniels. In search of the severity dimension of traffic events: Extended Delta-V as a traffic conflict indicator. Accident Analysis & Prevention. 2017 Jan 31;98:46-56

[Lee76] D.N. Lee. A theory of visual control of braking based on information about time-to-collision. Perception. 1976 Dec; 5(4):437-59

[MFH17] S.S. Mahmud, L. Ferreira, M.S. Hoque, and A. Tavassoli. Application of proximal surrogate indicators for safety evaluation: A review of recent developments and research needs. IATSS Research. 2017

[O18a] K. Czarnecki. Operational World Model Ontology for Automated Driving Systems – Part 1: Road Structure. Waterloo Intelligent Systems Engineering Lab (WISE) Report, University of Waterloo, 2018

[O18b] K. Czarnecki. Operational World Model Ontology for Automated Driving Systems – Part 2: Road Users, Animals, Other Obstacles, and Environmental Conditions. Waterloo Intelligent Systems Engineering Lab (WISE) Report, University of Waterloo, 2018

[Sch14] C. Schwarz. On computing time-to-collision for automation scenarios. Transportation research part F: traffic psychology and behaviour. 2014 Nov 30;27:283-94

[ST06] P. Songchitruksa, A.P. Tarko. The extreme value theory approach to safety estimation. Accident Analysis & Prevention. 2006 Jul 31;38(4):811-22

[Tar12] A.P. Tarko AP. Use of crash surrogates and exceedance statistics to estimate road safety. Accident Analysis & Prevention. 2012 Mar 31;45:230-40

[TSF16] National Highway Traffic Safety Administration (NHTSA). 2015 Motor Vehicle Crashes: Overview. Research note, DOT HS 812 318, 2016

[WJ12] K.F. Wu amd P.P Jovanis. Crashes and crash-surrogate events: Exploratory modeling with naturalistic driving data. Accident Analysis & Prevention. 2012 Mar 31;45:507-16

[YPA08] X. Ye, R.M. Pendyala, F.S. Al-Rukaibi, K. Konduri. Joint Model of Accident Type and Severity for Two-Vehicle Crashes. In Transportation Research Board 87th Annual Meeting 2008 (No. 08-2323)

[ZIM14] L. Zheng, K. Ismail, and X. Meng. Freeway safety estimation using extreme value theory approaches: A comparative study. Accident Analysis & Prevention. 2014 Jan 31;62:32-41