

A Core Ontology of Safety Risk Concepts Reconciling Scientific Literature with Standards for Automotive and Railway

Hermann Kaindl¹(✉), Thomas Rathfux¹, Bernhard Hulin², Roland Beckert¹,
Edin Arnautovic¹, and Roman Popp¹

¹ Institute of Computer Technology, TU Wien, Vienna, Austria
kaindl@ict.tuwien.ac.at

² Berner & Mattner Systemtechnik GmbH, Munich, Germany

Abstract. Safety is a major concern for both automobiles and railway vehicles. The related standards provide definitions of the same concepts such as *Risk*, *Harm*, *Hazard*, etc., which we consider here as the core concepts. However, related conceptual models existing in the scientific literature either are inconsistent or do not cover the core concepts comprehensively.

We modeled the core of these safety concepts ourselves both in meetings and with tool support, based on the definitions given in the related standards. As a result, this paper presents a small core ontology of safety risk concepts for reconciling the scientific literature with standards. Since it matches the terminology of the related standards, it may serve as a reference model in the future. In fact, we already used it ourselves for systematically studying where human error may compromise safety.

1 Introduction

In the context of our overall effort to support reuse in safety risk analysis (see, e.g., [18], we have been working on tool support. Such a tool needs to allow for input, handling and storing information on concepts like *Hazard*. Our chosen approach to generate parts of such a tool using Eclipse, a related *metamodel* has to be defined.

So, we looked up standards and related scientific literature to gather information for such a metamodel. Unfortunately, we found inconsistencies between the terminology of the standards with conceptual models in the literature. In addition, we could not find any conceptual model in the literature that would cover the core concepts comprehensively.

In particular, we investigated this issue in the context of automobiles and railway vehicles. In general, all such vehicles are covered by the generic standard IEC 61508 [3], which has the scope of Electrical / Electronical / Programmable Electronic Safety-related Systems (E/E/PE) and is based on ISO/IEC Guide 51 [5]. For practical reasons, more specific standards apply:




Fig. 1. Taxonomy of standards under investigation

- ISO 26262 [4] for automobiles, and
- EN 50126 [1] & TR 50126-2 [2] for railway systems.

Figure 1 depicts the relationships between these standards in the notation of the Unified Modeling Language (UML) [22], see omg.org for the current version. The arrow head points from the specific standards to the more general one (IEC 61508), which is also associated with ISO/IEC Guide 51.

ISO/IEC Guide 51:2014 [5] provides requirements and recommendations for the drafters of standards for the inclusion of safety aspects. This standard is applicable to any safety aspect related to people, property or the environment, or to a combination of these.

ISO 26262 is the functional safety standard for road vehicles and is derived from the generic functional safety standard IEC 61508. It deals with the possible hazards that could result from function failure in the electrical/electronic system in automotive vehicles.

EN 50126 is relevant for the whole railway system and not limited to railway vehicles. In contrast to IEC 61508 and ISO 26262, the railway standard EN 50126 is not limited to hazards resulting from malfunctioning of E/E/PE.

Primarily based on the terminology of these standards as defined in their glossaries, we started modeling of what we consider the core safety concepts. In addition, we employed tool support for finding relations between these concepts. In the course of several iterations over model versions in meetings, the models were most importantly extended and refined by expert knowledge from both the automotive and railway domains. We present here the resulting core ontology.

While it may have various applications in practice as a reference model in the future, we already used it for a preliminary but systematic study on where *human error* may compromise safety. This involved both traversing the graph of ontology concepts and looking at it as a whole.

The remainder of this paper is organized in the following manner. First, we motivate our work explicitly and discuss related work in the scientific literature. Then we elaborate on our effort on conceptual modeling of terminology from standards. Based on that, we explain our resulting core ontology of risk concepts. In addition, we sketch its fit into an *upper ontology*. As a possibility to make use of our ontology, we sketch how human error may compromise safety.

2 Motivation

While we originally strived for a metamodel for our tool support, our motivation for creating such a core ontology soon became more fundamental. In fact, safety assessment is in many ways subjective, partly because of individual risk perception, experience, education, cultural pressure and habits.

To reduce the arbitrariness of safety assessment, experts defined safety concepts such as *Risk*, *Hazard* and *Accident*. However, definitions in natural language are inherently ambiguous. With an ontology, at least the relations among the concepts contained can be made precise. They can also be visualized in figures as shown below, and such figures can support a common understanding of safety concepts.

In addition, even the definitions of safety core concepts such as *Risk* are not consistent between different safety standards. For example, while it is defined in the ISO 26262 standard [4] for the automotive domain as “combination of the probability of occurrence of harm and the severity of that harm”, for the railway domain it is defined as “the rate of occurrence of accidents and incidents resulting in harm (caused by a hazard) and the degree of severity of that harm” [2]. For creating our core ontology, we made an ontological decision in favor of the former definition, since it is actually derived from the ISO/IEC Guide 51 [5].

Being precise and consistent in this regard is actually a major concern in practice. This was, for instance, a major lesson learned by the author of this paper who is a safety expert in the railway domain, in an international project for the installation of a people mover. According to this real-world experience, if understandable and consistent definitions are not introduced in an early project phase, later much time will be wasted with discussions and with the reformulation of documents for the safety case.

Moreover, the consistency of definitions of safety concepts may become important after an accident in legal courts. Interpretations of safety concepts may be discussed there and related questions raised, such as the following:

- What was the interpretation of the safety concept x for the safety case?
- What are other interpretations of this safety concept (in the standards used or other similar ones)?
- Would the other interpretation have led to additional safety requirements?
- Could the accident have been avoided if such additional safety requirements were taken into account?

Our core ontology of safety risk concepts may help to answer such questions consistently.

3 Related Work in the Literature

Ambiguity of safety standard terminology and the problems resulting are discussed in [12, 23]. Models of safety standards can contribute to avoid misunderstandings and conflicting views on the concepts behind the terminology.

Such a model for IEC 61508, with the focus on creating a chain of evidence for safety compliance demonstration, is proposed in [21]. Unfortunately, as explained below, there are ontological problems with this model, in particular its *Risk* concept. Another model of a few concepts from ISO/IEC Guide 51, from which IEC 61508 takes many of the core glossary definitions, can be found in [23]. It only centers around a model of *Risk*, but also this model has ontological problems as explained below. Hence, we could not base our core ontology on either of these papers.

A discussion of evolving definitions of the concepts *Risk*, *Hazard* and *Mishap* in military standards is discussed in [28]. As a result, a formalized model for calculating hazard and mishap occurrence probabilities is presented. The ontological view of risk-related concepts in these military standards is quite different from the one in automotive and railway standards. Hence, it was not possible to base our core ontology on this work, either. However, the increasing importance of the *Mishap* concept in the evolution of military standards suggests to us the importance of the related *Accident* concept. As explained below, the inclusion of *Accident* into our core ontology was only in the course of an evolution of our conceptual models.

The closest attempt to our ontological modeling in this paper can be found in [13], where we focused on the differences in automotive and railway standards and unified them conceptually as far as possible. In contrast, the current paper provides a core ontology of the common safety concepts. In addition, there was no model of *Risk* yet in [13].

Our development of a core ontology of safety risk concepts may be considered as a simple application of *Ontology Engineering* (OE) [8]. OE represents “the set of activities that concern the ontology development process, the ontology life cycle, and the methodologies, tools and languages for building ontologies” ([9], as cited in [25]. Typical activities in OE are Domain Analysis and Specification (knowledge acquisition, and the definition of ontological purpose, including its use cases, users, etc.), Conceptualization (structuring of domain knowledge), and Implementation (expressing the ontology using an appropriate ontology representation language). On top of the activities for ontology building are the activities for ontology utilization and application (e.g., building tools for the defined use cases). Another important activity in OE is ontology *evaluation*. The goal of ontology evaluation is to estimate the quality of the ontology, and it includes ontology *validation* (investigation if the ontology represents the real-world domain concepts and their relationships appropriately, and if it fulfills the ontology use case and purpose), and *verification* (proving consistency and that the ontology is correctly constructed according to the language used, etc.) [10]. However, most of the ontology evaluation approaches [6] deal with large, complex and more or less formally represented ontologies (e.g., in OWL or description logic) and are not suitable for our case.

Since we have used a semi-formal representation in UML without constraints, logic formalisms, etc., and having in mind the current size of our core ontology, explicit ontology verification is not feasible. Regarding validation, we iteratively

reviewed the results from ontology development steps using expert knowledge. In our future work, we plan to validate the ontology against use cases. Another option for ontology validation would be to automatically create an ontology from standards and (or) scientific literature and qualitatively compare it to our manually created ontology. Sfar et al. [24] use a similar comparison to evaluate automatically created ontologies against a “gold-standard” ontology created by humans. So, using ontology learning [27] for the validation of our core ontology would be a valid goal for our future research. We already gained first experience in comparing semi-automatically created taxonomies (light-weight ontologies) to manually created domain models in requirements engineering [7].

4 Conceptual Modeling of Terminology from Standards

Conceptual modeling is, in general, not that simple. Regarding models of safety concepts in the literature, we particularly found inconsistencies in [23] and in [21]. In both cases, these are supposedly related to misunderstandings of the *aggregation* relationship of UML.

In [23], a categorization of the concept *Risk* is correctly modeled using *generalization* of the classes representing the subconcepts in UML. However, “Damage” (supposedly used here as a synonym of *Harm*) is modeled there as an aggregation of three special cases of *Harm*, and this should rather be modeled as well using *generalization*.

In [21], the concept *Risk* is modeled as class with a few UML *attributes*, including “likelihood” and “consequence”. Assuming that they correspond to *Probability* and *Severity* according to the standards that we model below, there is an interesting modeling issue. In the specification of UML, an attribute is said to be “semantically equivalent to a composition association”. When considering this statement more precisely, the question arises, in which sense an attribute is part of an object. In the UML *metamodel*, *attribute* is part of *class* in a composition. In this sense, an attribute is an entity of its own, which defines UML. But in the specification of UML as well as in [22], attributes are also said to be “composition relationships between a class and the classes of its attributes”. In this sense, an attribute would model the same relationship as a composition. A simple example shows that this view is questionable. The region of a wine can be modeled as its attribute (as one of possibly several), but this does not mean that any particular region is “part of” a particular wine. Already in [26], “attribution” was said to be often confused with a whole-part relationship. The argument that these are different relationships was another simple example: “While towers have height as one of their attributes, height is not a part of a tower.”

Therefore, it is rather the class representing the concept *Harm* that may have (among others) the attributes *Probability* and *Severity*, see Fig. 2. While we think that this is a ‘true’ model of these concepts (according to the standards under consideration), this way of modeling raises yet another issue. How would it be possible in such a model to represent that this combination of *Probability* and *Severity* of *Harm* is *Risk*? All this justifies the ontological decision to model this inner core as given below (using aggregation).




Fig. 2. Harm class with attributes

For the actual modeling involved for achieving our proposed core ontology, we pursued two different ways. We employed tool support for getting automated suggestions for association relationships, and we had a series of expert meetings, i.e., meetings involving two safety experts. Note, that the tool run was only after the second of a total of four meetings. So, it was not intended to bias the whole effort but only to see more exactly what can be extracted directly from the given glossaries.

4.1 Tool-Supported Modeling

After the second meeting (as sketched below), we tried tool-supported modeling. We were interested in getting suggestions for (binary) association relations between any two of the core concepts under discussion in the meetings, based on their glossary definitions in the standards under investigation. Our major interest was to see what exactly these definitions say about potential relations between the concepts defined.

More precisely, we employed the tool RETH (Requirements Engineering Through Hypertext), a tool for requirements specification according to the method with the same name. RETH combines object-oriented technology and hypertext. It was developed under the guidance of the first author of this paper some time ago, see, e.g., [16].

For tool-supported modeling in the course of creating our core ontology of safety risk concepts, we used the RETH tool to automatically generate *glossary links*, see [17]. More generally, it is a semi-automated generation that allows the user to reject a suggested link, but we refrained from this option in order not to influence the result. Based on such links, we let the tool automatically generate (binary) association relations in a second step, see [15]. According to the heuristic behind that, RETH simply generates an association, if and only if there is a glossary link in either direction. Of course, such proposed associations can be deleted manually, e.g., if they are transitive and, therefore, may be considered redundant. Again, we refrained from this option in order not to influence the result. Note, that this tool can also propose generalizations, e.g., for *Risk* being more general than *Individual Risk* (based on an obvious linguistic clue), but we did not have such a case here.

Let us show an example of an entry for a concept and its definition as an excerpt from a linearized tool output:

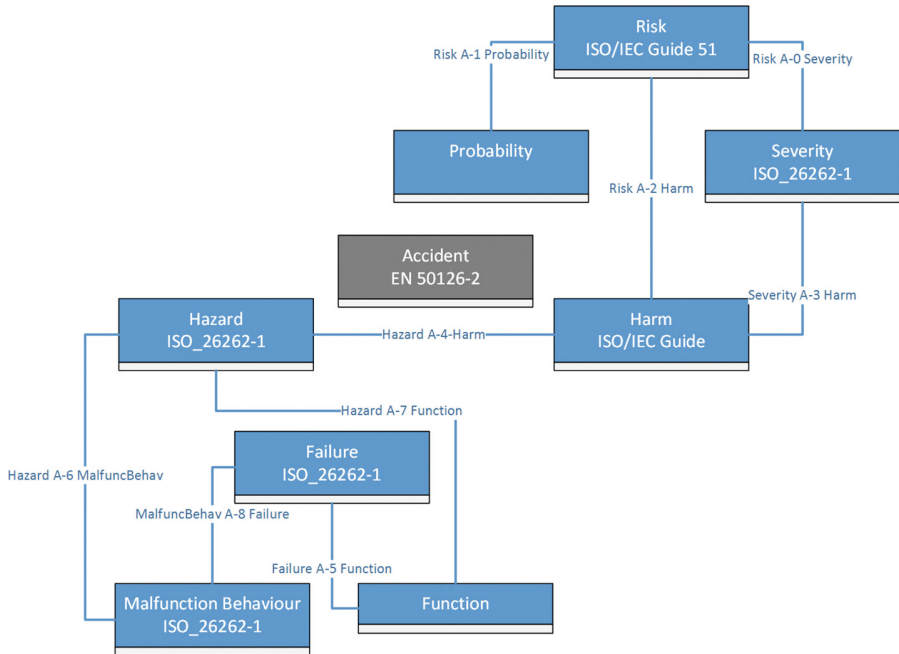


Fig. 3. Conceptual model with tool-generated associations

Risk

– Source of Definition

ISO/IEC Guide51:2014

- combination of the probability of occurrence of harm and the severity of that harm
- **A-0 Severity**
- **A-1 Probability**
- **A-2 Harm**

The link to “harm” was already given in this standard, but not the one to “severity”, which was generated by the tool. The associations in this case only correspond to out-going links from this concept and are shown here through a generated name and a link to the associated concept.

The resulting model from the tool run is shown in UML in Fig. 3. With respect to our inner core of concepts around the concept *Risk*, the association A-2 is a typical case of a redundant transitive relation, which can be deleted in order not to clutter the diagram. A-0 and A-1 are simply shown here as associations, while they may be modeled as their special case of an aggregation in UML. However, the tool does not have any clue for such a distinction. (Note, that the UML definition of an aggregation is vague, and attempts to formalize them in logic are difficult.)

An interesting observation is that the concept *Accident* is shown here in isolation, i.e., without any association relation with any of the other concepts of this model. Hence, let us have a look at its definition:

Accident

– Source of Definition

EN 50126-1

- an unintended event or series of events resulting in loss of human health or life, damage to property or environmental damage

In fact, there is no link that could have been found by the tool, while this text can be interpreted in such a way that *Accident* may be related to *Harm*. This concept is directly associated in the model with the concept *Hazard*, however, based on the following definition:

Hazard

– Source of Definition

ISO 26262-1

- potential source of harm caused by malfunctioning behaviour of the item
- **A-4 Harm**
- **A-6 Malfunctioning Behaviour**
- **A-7 Function**

The reader is encouraged to compare this model with the ones created and elaborated at the meetings as sketched below, especially regarding this direct association. Note, in addition, that the concept *Malfunctioning Behaviour* was finally not included into our core ontology, although it would make sense, but it seemed to be less important in the standards under investigation.

4.2 Expert Meetings

As indicated above, we had four expert meetings including two safety experts, one primarily in the railway domain, the other in automotive. The other participants have primarily background in software and symbolic modeling, in particular also on ontologies. Note, that all participants of these meetings are also authors of this paper. Each meeting had five to six participants, and the duration was, on average, approximately seven hours. Between these meetings, we aligned ourselves via email and telecommunication, while we primarily worked on different tasks.

The starting point was a metamodel intended to create an Eclipse-based tool for supporting reuse of safety risk analyses. This metamodel included among other classes for requirements, etc., the following ones: *Function*, *Failure*, *Hazard*, *Severity* and *Tolerable Hazard Rate* (see also [18]).

Since this metamodel was considered insufficient by these authors, the relevant standards were consulted, first IEC 61508 and ISO Guide 51. Since the




Fig. 4. White board with sticky notes from second meeting

terms in these standards are not unambiguously defined, we decided to look for conceptual models that we could adopt for the metamodel needed. Unfortunately, as explained above, we could not find a comprehensive model of the core safety terms as needed in the scientific literature. It even contained conceptual models that are inconsistent with the terminology of these standards.

In the course of a first meeting of all the authors, we primarily discussed an integration of a *Risk* model with the concepts corresponding to the classes of our previous metamodel. Immediately after this first meeting, however, the safety expert of the railway domain criticized that *Accident* was missing in our model and pointed to the definition according to EN 50126. Additionally, he proposed to introduce *Hazardous Situation* for representing preconditions that could lead to an *Accident*.

In our second meeting, we discussed possible inclusions of these concepts into our model. Even though neither IEC 61508 nor ISO 26262 define *Accident* explicitly, we decided to extend our model with this concept. In order to determine reasonable associations between the given concepts, we used sticky notes with a concept name per note, and arranged them on a white board (for the result see Fig. 4). As shown above, an association between *Accident* and *Harm* is obvious, but what causes the occurrence of such an unintended event? The definitions do not clarify that. So, the safety experts' knowledge was brought in and

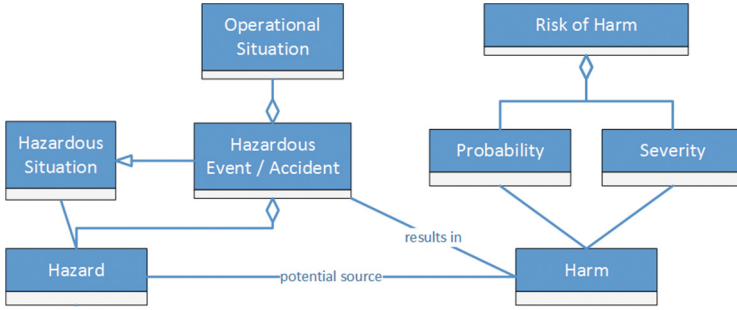


Fig. 5. Conceptual model in the course of the third meeting

led to the inclusion of *Accident* between *Hazard* and *Harm*. Still, we could not determine associations of *Hazardous Situation* with the other concepts, although some relation with *Hazard* is suggestive.

Between the second and the third meetings, we used the RETH tool as explained above. While the resulting model is fairly similar to the one after our second meeting, the tool could not find any association of *Accident* with the other concepts.

In the third meeting, we examined the association between *Hazard* and *Accident* more closely. In particular, we took ISO 26262 into account, which defines *Hazardous Event* as the combination of a *Hazard* and an *Operational Situation*. According to the glossary definitions, there is a missing link between *Hazard* and *Accident*, because an *Accident* is a result of a single event or a series of events. After long discussion, we erroneously decided to add both the concepts *Hazardous Event* and *Accident* as one named *Hazardous Event / Accident* due to their apparent similarity, resulting in the conceptual model given in Fig. 5.

After even more discussion in the course of the third meeting, we split *Hazardous Event* and *Accident*, and defined an association named “may cause” between them.

In the fourth meeting, we reviewed the resulting model from the third meeting and did not find a flaw, while there are always options for other ontological decisions. The only change was adding *Triggering Event*, with an association named “triggers” with *Accident*. This model intends to reflect that both *Hazardous Event* and *Triggering Event* are preconditions of an *Accident*.

5 Our Core Ontology

The resulting conceptual model is shown in Fig. 6. We consider it as a core ontology of safety risk concepts. While our sketch of its evolution above should already serve as an explanation, a few more explanations are still necessary for the rationale of some of its parts.

The ontological decision for the aggregation of *Risk* (instead of attributes) is explained above in detail. Such an aggregation relation in UML is shown as a




Fig. 6. Resulting conceptual model defining a core ontology of safety risk concepts

diamond, see Figs. 5 and 6. The rationale for the aggregation of *Hazardous Event* is by analogy. In fact, both underlying definitions in the standards use the same formulation “combination of”.

The name of the association between *Accident* and *Harm*, “results in”, suggests that every accident results in harm. Otherwise, this unintended event or series of events would not be considered an accident according to the definition of *Accident* in EN 50126 (see also above). Instead of naming it “may results in”, which somehow involved yet another probability, our railway expert suggested to assign *Severity 0* in case there is no resulting human *Harm*. In this way, our model resolves the very narrow and conflicting definition of *Harm* in the ISO 26262 standard, which is restricted to human health but not to goods or the environment.

The concepts *Function* and *Failure* are only relevant for ‘Functional Safety’, i.e., when assessing *Harm* based on analyzing potential failures of each function of a system. This has to be done according to ISO 26262 in the automotive domain. For other kinds of safety analyses, these concepts may be ignored.

Overall, this core ontology is an interesting combination of both the automotive and the railway domains. The concepts shown in blue are common, while the others are based on terminology from ISO 26262 and EN 50126, respectively.

6 Upper Ontology

For such an ontology, also its fit into a so-called *upper ontology* is important. Upper ontologies represent general concepts to be used for creating more specific domain ontologies such as ours. In this regard, let us focus on a specific problem.

In our meetings, we struggled with the terms *Hazardous Situation* and *Hazardous Event*. In fact, there is no clear conceptual differentiation between the concepts *Situation* and *Event* in the standards that we used. To clarify this, we looked into several upper ontologies. OpenCyc [19] is the largest and the best known upper ontology, containing around 10^5 generic concepts. OpenCyc defines the concept of a *Situation* roughly as a state and as specializations of *Intangible* and *Temporal* concepts. *Event* is defined as a specialization of *Situation* (a dynamic situation in which the state of the world changes). Contrary to *Event*, in a *Static Situation* (as another specialization of the concept *Situation*), objects and their relations do not change over time. OpenCyc also defines the concept of *DangerousSituation* as a specialization of *Situation* where “a significant risk of death, injury, or property damage exists”. Even though these definitions are not precise enough for safety analysis (e.g., *DangerousSituation* has the word “Hazard” as a synonym, but the safety standards explicitly distinguish between these two concepts), it seems as though the conceptualizations of OpenCyc fit well the intrinsic meaning of the safety-related concepts of our ontology.

Deeper investigation on the relation to other upper ontologies (e.g., ABC upper ontology [14], General Formal Ontology [11], or SUMO — Suggested Upper Merged Ontology [20]) will be part of our future work.

7 Human Error

Human error may compromise safety in the context of safety-critical systems. Let us sketch how our core ontology of safety risk concepts can be used for a systematic analysis of human error.

First, the graph of ontology concepts can be traversed systematically, where each concept is investigated regarding human error. In particular, additional hazards caused by human error are important. They may lead to additional hazardous events and accidents as well. With respect to user errors, e.g., especially the triggering events of accidents are of interest. Design and development errors as well as manufacturing, construction and installation errors seem to be more related to functions and related failures. Also maintenance errors are to be studied in this context.

With the help of the ontology, human error can be identified or classified systematically, see Table 1. The three concepts *Function*, *Failure* and *Hazard* at the left of Fig. 6 are relevant for human error analysis if an operator or maintainer is involved in fulfilling or supporting a function, respectively. For example, a train driver of certain railway vehicles has to fulfill part of the so-called *Parking* function. For the analysis of human error in such a case, the acting humans are considered part of the system. Such an analysis is especially important in degraded modes with more intensive use of human capabilities. In general, the traversal of the ontology guides from each instance of *Function* to analyzing it with regard to *Failure* and *Hazard*. This functional safety analysis has to investigate which failures may be caused by such a human error, and which hazards may result. Analogously to operator error, errors of maintainers can be analyzed

in this way. For example, in a railway vehicle the generation of pressurized air is safety-relevant, since many functions such as braking are implemented based on it. Filtering of dust and dirt in the pressurized air is a crucial point, since air pipes may become locked by dust and dirt. Hence, the dust filters have to be changed after at least one year. If this does not happen, e.g., caused by human error in maintenance, these filters will lose their required function.

Table 1. Classification of human error related to concepts of the ontology.

Role of human	Related concepts	Possible reason
Operator	<i>Function, Failure, Hazard</i>	Operator involved in <i>Function</i>
Maintainer	<i>Function, Failure, Hazard</i>	Maintainer supporting <i>Function</i>
Person at risk	<i>Hazardous Event, Accident</i>	Self rescue
User	<i>Triggering Event, Accident</i>	Misuse

Affected persons of operator errors can, of course, be the operator who caused the hazard, but also others, such as passengers of a train. Related to persons at risk, i.e., all involved humans that may suffer *Harm* from a given *Hazard*, the concepts *Hazardous Event* and *Accident* are particularly important for the analysis of human error. Their instances need to be analyzed especially regarding possibilities for escaping or avoiding any instances of *Harm*, e.g., through self rescue, and what kind of human error may happen in this course. The avoidance of harm must be recognizable, understandable, possible and desirable. For example, in case of large and abnormal vibrations in a wagon of a train, pulling the emergency brake may be the most appropriate action to be taken by a passenger. However, there are some problems involved in such a situation. First of all, the passenger needs to recognize that these vibrations are abnormal. Given that, the passenger has to understand that it is reasonable or necessary for him or her to act. In addition, the passenger needs to figure out which actions are possible, e.g., pulling the emergency brake or moving to another part of the train. Finally, the passenger needs to judge that such an action is desirable, since unjustified pulling an emergency brake is also subject to being punished, and decide to actually perform such an action. In particular, such an analysis of human error needs to take into account that humans involved in an accident are usually under stress, and the more stressed humans are the more likely they commit errors.

The concept of a *Triggering Event* related to an *Accident* is relevant in the context of unintended or intended misuse. An example of a triggering event is pushing the button for opening doors of a train during the *Operational Situation* in a tunnel at high speed (say, 300 km/h). In such a situation, the aerodynamic forces can be strong enough to pull a passenger out of the train if a door in the vicinity opens. This human error of misuse is covered by electronic locking of the doors, where unlocking a door in such a situation would be an instance of *Hazard*.

Such a systematic analysis of human error may, in turn, suggest an integration of additional technical assistance systems. These are intended to reduce the possibilities of human error or its negative effects. The overall safety assessment needs to find a balance between human and technical aspects related to hazards and risks.

Another potential use of our core ontology related to human error is to look at it as a whole. After all, it is currently fed into a tool for supporting reuse of risk analyses (through a related metamodel). When risk analyses with all the related information according to our model will be reused for similar cases, e.g., previous hazards will be taken into account that otherwise may be overlooked by human error.

Even regarding standards, both their creation and their application, there is some potential use of our core ontology. After all, it is based on ISO 26262, EN 50126 and IEC 61508. Problems often arise from contradicting or arbitrary definitions, or even missing definitions. For example, in ISO 26262 the term “accident” is not defined even though it is used in some of its parts. For the creation of future (versions of) standards, human error may be reduced through this and enhanced ontologies.

8 Conclusion and Future Work

Primarily based on the glossaries of standards for automotive and railway, we created conceptual models, both using tool support and in a series of expert meetings. Especially in these meetings, of course, expertise of two safety experts played a major role in the evolution of these models. We consider the resulting model a core ontology of safety risk concepts covering both domains, which also fits into a major upper ontology. As a preliminary application of this core ontology, we used it for systematically studying possibilities of human error compromising safety.

In on-going and future work, we base a corresponding metamodel for tool creation using Eclipse on this core ontology. This metamodel will also include requirements-related concepts, which we have sufficient previous experience with. Using the resulting tool, and indirectly our core ontology, we will perform case studies, focusing on reuse of risk analyses. In this course, we will particularly investigate whether this reuse can help to reduce human error of omitting important information on previously known hazards, etc. All this will be important for the sake of validation of our proposed approach. Also extending the scope, e.g., to avionics or healthcare will be of interest.

Acknowledgment. Part of this research has been carried out in the RiskOpt project (No. 845610), funded by the Austrian BMVIT (represented by the Austrian FFG).

References

1. EN 50126-1, Railway applications – The specification and demonstration of reliability, availability, maintainability and safety (RAMS). Part 1: Basic requirements and generic process, September 1999

2. CLC/TR 50126-2, Railway applications – The specification and demonstration of reliability, availability, maintainability and safety (RAMS). Part 2: Guide to the application of EN 50126-1 for safety, February 2007
3. IEC 61508, Functional safety of electrical/electronic/programmable electronic safety-related systems, May 2010
4. ISO 26262, Road vehicles - Functional safety, November 2011
5. ISO/IEC Guide 51 - Safety aspects - Guidelines for their inclusion in standards (2014)
6. Brank, J., Grobelnik, M., Mladenčić, D.: A survey of ontology evaluation techniques. In: Proceedings of 8th International Multi-conference Information Society, pp. 166–169 (2005)
7. Casagrande, E., Arnautovic, E., Woon, W.L., Zeineldin, H.H., Svetinovic, D.: Semi-automatic system domain data analysis: a smart grid feasibility case study. *IEEE Trans. Syst. Man Cybern. Syst.* **PP**(99), 1–11 (2016)
8. Fernandez-Lopez, M., Gomez-Perez, A., Juristo, N.: Methontology: from ontological art towards ontological engineering. In: Proceedings of the AAAI 1997 Spring Symposium, Stanford, USA, pp. 33–40, March 1997
9. Gómez-Pérez, A., Fernández-López, M., Corcho, O.: *Ontological Engineering: With Examples from the Areas of Knowledge Management, e-Commerce and the Semantic Web*. Advanced Information and Knowledge Processing. Springer, New York Inc. (2007)
10. Hajdu, M., Skibniewski, M.J., Bilgin, G., Dikmen, I., Birgonul, M.T.: Selected papers from creative construction conference 2014 ontology evaluation: an example of delay analysis. *Procedia Eng.* **85**, 61–68 (2014). <http://www.sciencedirect.com/science/article/pii/S1877705814018955>
11. Herre, H.: General formal ontology (gfo): a foundational ontology for conceptual modelling. In: Poli, R., Obrst, L. (eds.) *Theory and Applications of Ontology*, vol. 2. Springer, Berlin (2010)
12. Hogganvik, I., Stolen, K.: Risk analysis terminology for IT-systems: does it match intuition? In: 2005 International Symposium on Empirical Software Engineering 2005, p. 10, November 2005. <http://dx.doi.org/10.1109/ISESE.2005.1541810>
13. Hulin, B., Kaindl, H., Rathfux, T., Popp, R., Arnautovic, E., Beckert, R.: Towards a common safety ontology for automobiles and railway vehicles. In: *European Dependable Computing Conference* (2016, to appear)
14. Hunter, J.: Enhancing the semantic interoperability of multimedia through a core ontology. *IEEE Trans. Circ. Syst. Video Technol.* **13**(1), 49–58 (2003)
15. Kaindl, H.: How to identify binary relations for domain models. In: *Proceedings of the Eighteenth International Conference on Software Engineering (ICSE-18)*, pp. 28–36. IEEE, Berlin, March 1996
16. Kaindl, H.: A practical approach to combining requirements definition and object-oriented analysis. *Ann. Softw. Eng.* **3**, 319–343 (1997)
17. Kaindl, H., Kramer, S., Diallo, P.S.N.: Semiautomatic generation of glossary links: a practical solution. In: *Proceedings of the Tenth ACM Conference on Hypertext and Hypermedia (Hypertext 1999)*, Darmstadt, Germany, pp. 3–12, February 1999
18. Kaindl, H., Popp, R., Raneburger, D.: Towards reuse in safety risk analysis based on product line requirements. In: *2015 IEEE 23rd International Requirements Engineering Conference (RE)*, pp. 241–246, August 2015
19. Lenat, D.B.: Cyc: a large-scale investment in knowledge infrastructure. *Commun. ACM* **38**(11), 33–38 (1995). <http://doi.acm.org/10.1145/219717.219745>

20. Niles, I., Pease, A.: Towards a standard upper ontology. In: Proceedings of the International Conference on Formal Ontology in Information Systems, FOIS 2001, vol. 2001, pp. 2–9. ACM, New York (2001). <http://doi.acm.org/10.1145/505168.505170>
21. Panesar-Walawege, R.K., Sabetzadeh, M., Briand, L., Coq, T.: Characterizing the chain of evidence for software safety cases: conceptual model based on the IEC 61508 standard. In: 2010 Third International Conference on Software Testing, Verification and Validation, pp. 335–344, April 2010. <http://dx.doi.org/10.1109/ICST.2010.12>
22. Rumbaugh, J., Jacobson, I., Booch, G.: The Unified Modeling Language Reference Manual. Addison-Wesley, Reading (1999)
23. Schnieder, L., Schnieder, E., Ständer, T.: Railway safety and security —two sides of the same coin? In: International Railway Safety Conference 2009 (2009). <http://www.intlrailsafety.com/bastad/20090928/09-stander/paper.pdf>
24. Sfar, H., Chaibi, A.H., Bouzeghoub, A., Ghezala, H.B.: Gold standard based evaluation of ontology learning techniques. In: Proceedings of the 31st Annual ACM Symposium on Applied Computing, SAC 2016, pp. 339–346. ACM, New York (2016). <http://doi.acm.org/10.1145/2851613.2851843>
25. Simperl, E.P.B., Tempich, C.: Ontology engineering: a reality check. In: Meersman, R., Tari, Z. (eds.) OTM 2006. LNCS, vol. 4275, pp. 836–854. Springer, Heidelberg (2006). http://dx.doi.org/10.1007/11914853_51
26. Winston, M.E., Chaffin, R., Herrmann, D.: A taxonomy of part-whole relations. *Cogn. Sci.* **11**, 417–444 (1987)
27. Wong, W., Liu, W., Bennamoun, M.: Ontology learning from text: a look back and into the future. *ACM Comput. Surv.* **44**(4), 20:1–20:36 (2012). <http://doi.acm.org/10.1145/2333112.2333115>
28. Zhao, N., Zhao, T.: An event-chain risk assessment model based on definition evolution in safety criterions. In: 2011 9th International Conference on Reliability, Maintainability and Safety (ICRMS), pp. 573–578, June 2011. <http://dx.doi.org/10.1109/ICRMS.2011.5979333>