

A Modification to the Hazardous Situation ODP to Support Risk Assessment and Mitigation

Michelle Cheatham¹, Holly Ferguson², Charles Vardeman II², and Cogan Shimizu¹

¹ Wright State University

{michelle.cheatham,cogan.shimizu}@wright.edu

² University of Notre Dame

{hfergus2,cvardema}@nd.edu

Abstract. The Hazardous Situation ontology design pattern models the consequences of exposure of an object to a hazard. In its current form, the ODP is well suited for representing the consequences of exposure after the fact, which is very useful for applications such as damage assessment and recovery planning. In this work, we present a modification to this pattern that enables it to additionally support proactive questions central to risk assessment and mitigation planning.

Keywords: hazard, ontology design pattern, risk assessment, risk mitigation

1 Introduction

As defined in [1], a hazard is “a *potential* source of harm to someone or something.” The concept of a hazard is central to answering very important questions in a variety of domains. For example, assessment of an individual’s finances prior to approving or denying a loan request relies on the ability to enumerate and explore the details of the financial hazards that may impact the individual. Developing architectural and construction plans for a new building necessitates consideration of hazards such as earthquakes, fires, and high winds in order to mitigate the risks presented by each. And establishing safety protocols and incident response plans at a chemical plant requires an understanding of the chemicals involved and their potential consequences when touched, inhaled, or otherwise interacted with.

Ten years ago, the U.S. National Science Foundation funded a committee to assess the current state of research related to understanding the “societal responses to natural, technological, and willful threats.” The summary of the committee’s report [6] laments that hazard and disaster policy is often reactive rather than proactive and points out that (proactive) risk mitigation and (reactive) disaster management should not be treated as isolated fields. Additionally, the report states that lack of data accessibility is hindering progress in both of these fields. As we will see in Section 2, there has been significant work on modeling various aspects of risks, risk management, and incident response. However,

existing work often takes an implicitly reactive approach by representing the consequences of hazards after they have been involved in an event. In this work, we seek to model exposure to a hazard in such a way that data organized according to this model can be used for *proactive* analysis. In particular, our goal is to arrive at an ontology design pattern that is capable of following competency questions:

- What is the consequence of exposing a particular object to a particular hazard for a particular amount of time?
- What are the hazards to a particular object at a particular geographic location?
- What hazards are capable of causing a given consequence?
- What susceptibilities could result in negative consequences to an object exposed to a particular hazard?
- What mitigation strategies could be used to effect exposure to a particular hazard?

2 Related Work

In this section we review previous efforts to model the concept of a hazard. While there have been many efforts by government agencies, insurance companies, and other organizations to develop taxonomies or other classification systems for hazards, this review focuses explicitly on ontology-based modeling approaches and considers these other efforts to be out of scope.

Some of the prior work in this area has focused on modeling hazards within a particular domain of interest. For example, Letia and Groza present a hazard ontology to support risk analysis in the food supply chain in [2]. The types of hazards modeled are constrained to biological (e.g. bacteria, parasites), chemical (e.g. mercury, unsafe preservatives), and physical (e.g. fish bones, plastic) hazards related to food that might cause illness or injury to a consumer. Malaal and his colleagues attempt something similar for the railroad transportation domain, though the ontology they present in [3] does not follow best practices.

Other work has taken a more domain-agnostic approach. One example is [8], in which an ontology is established to model the analysis done after a hazard has occurred (a so-called “post mortem”). Rather than focusing on the hazard itself, this ontology is meant to capture the complete hazard evaluation process (from identification, to evaluation, to fault analysis) and its conclusions. Other domain-agnostic work has focused more specifically on modeling hazards. Winther and Marsh point out that it is not always clear what element in a situation is the hazard [7]. They attempt to clarify this point by defining a hazard to be at the intersection between two subsystems; however, this approach has not been widely adopted and does not seem intuitive to domain experts. The only other approach to modeling hazards in a domain agnostic way of which we are aware is that done by Lawrynowicz and Lawniczak in [1]. The hazardous situation ODP considers a *hazardous event* to be an event in which an *object* is *exposed* to a *hazard*. A hazardous event can have a *cause* and a *consequence*. A *hazardous situation* is then defined as a situation that *participates in* a hazardous event.

All of these domain agnostic approaches implicitly focus on representing exposure to a hazard after it has occurred. For instance, John Doe was exposed to poison ivy and developed a rash. Such information has many important applications related to damage assessment and recovery operations. However, we are interested in modeling hazards in a way that also supports risk analysis and forward-looking mitigation planning. This goal places more requirements on the ODP. In particular, not every object will react to the same exposure to a hazard in the same way. For example, approximately 15 percent of people are not allergic to poison ivy. If Joe Smith is one such person, then we could represent his exposure to poison ivy using the same pattern, but with a different instance of the *consequence* entity. However, this does not allow us to answer prospective (as opposed to retrospective) questions along the lines of “Is it safe for Joe Smith to enter a particular hazardous situation?” Additionally, a critical component to proactively reducing risk is the concept of a mitigating factor. For example, if John Doe touched poison ivy but was wearing gloves at the time, then he would not experience the same consequence. In the hazardous situation ODP, for example, this can be handled through the *exposure* entity (i.e. Mr. Doe’s exposure to the poison ivy is less if he is wearing gloves). However, exploring the impact of various mitigation strategies in a proactive way, for instance by asking questions like “Is it safe for John Doe to enter a particular hazardous situation if he implements a given mitigation strategy?”, requires making this concept explicit in the pattern.

3 Design Considerations

The existing hazardous situation ODP is well designed and already supports many of the core competency questions commonly asked related to exposure to hazards [1]. We therefore do not seek to reinvent the wheel, but rather to break out some aspects of exposure to a hazard into constituent components, such that the modified pattern can be used for risk assessment and mitigation planning.

One key element necessary to achieving our goals is the ability to represent that different objects may be impacted differently by the same level of exposure to a hazard, based on their individual susceptibility to that hazard. This is somewhat similar to the concept of affordances, in which the same object may afford different possibilities for action to different individuals. For instance, Ortman and Kuhn explain in [4] that whether or not a set of steps affords climbing depends on the relation between the height of each step and the length of an individual’s legs. Whether or not an affordance is invoked (e.g. the person actually climbs the stairs) is at the person’s discretion, however, whereas whether or not exposure to a hazard to which a person is susceptible causes some consequence is generally outside of that person’s control (and indeed can occur whether or not the person is even aware of the hazard).

A closer analogy comes from the domain of biological dispositions. Röhl and Janssen’s excellent paper on the ontology of dispositions [5] defines them as “a causal property that is linked to a realization, i.e. to a specific behavior which the individual that bears the disposition will show under certain circumstances or as response to a certain stimulus (trigger).” They discuss representing dispositions

at both the level of classes (e.g. all aspirin has the disposition to relieve pain) and instances (e.g. John Doe has the disposition to bleed excessively). Their ontology involves four key entities: the *Material Entity* that possesses the disposition (e.g. John Doe), the *Disposition* itself (e.g. tendency to bleed excessively), the *Quality* of the material entity that is the base of the disposition (e.g. insufficient clotting factor), and the *Realization* of the disposition in response to some trigger (e.g. massive blood loss after a cut). We make use of these concepts in our modifications to the Hazardous Situation ODP, described below.

4 Formalization

In this section we present the modifications to the Hazardous Situation ODP (shown graphically in Figure 1), describe the relevant entities, and formalize applicable axioms. In an effort to make this discussion more clear, we will use a simple (entirely hypothetical) example of Michelle, who is lactose intolerant and yet eats some ice cream.

The yellow boxes in Figure 1 represent entities from the original Hazardous Situation ODP, while the green boxes represent new entities.³ In addition to the four new entities, we have made one other change: the range of the *involves* property is now *Exposure* rather than *Hazard*. This was done because an object is not susceptible to a hazard, but rather to some exposure (i.e. dosage or amount) of the hazard.

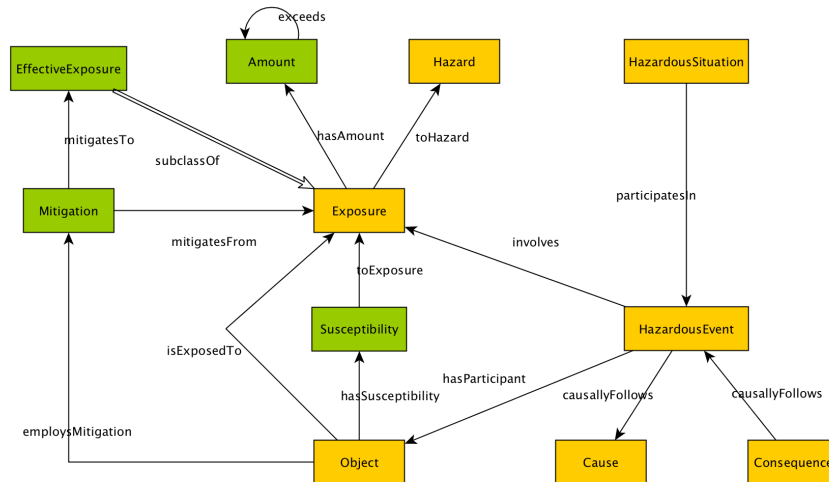


Fig. 1: The modified Hazardous Situation ODP

³ There is a *isSubclassOf* relationship between *Hazard* and *Object* that was omitted from the figure for clarity.

4.1 Entities

Object: Any physical entity. The object can be modeled at whatever level of granularity is necessary, e.g constituent parts of an overall object can be represented individual. This class is similar to a Material Entity in [5]. In our example, this is Michelle.

Exposure: The subjection of an object to a particular amount of a hazard. Similar to a Trigger in [5]. Here, the eating of the ice cream.

Hazard: Any entity capable of producing a (negative) consequence. This does not have to be material, or even physical. Examples include earthquake, fire, or nightmares. In this case, the hazard is ice cream.

Amount: The amount of a hazard to which an object is exposed. This is deliberately vague and will vary based on the type of hazard involved. For example, an amount of radiation might be 100 rem, while an “amount” of nightmares might be “terrifying.” In this example, the amount is two scoops.

Exceeds: An object property used to indicate when one amount exceeds another. This is used in the axioms to recognize when an object’s exposure to a hazard has exceeded its susceptibility and therefore triggers a corresponding consequence.

Susceptibility: The level of exposure to a hazard at which an object may experience a consequence of the hazard. Similar to Disposition in [5]. In this case, it is a relationship between Michelle and an exposure of half a scoop of ice cream.

Mitigation: Anything that changes either an object’s amount of exposure to a hazard or the exposure amount of an object’s susceptibility to that hazard. Examples include wearing a mask while around sick people (to reduce the number of germs you are exposed to) or getting a vaccination (to increase the number of germs you can be exposed to without becoming ill). In this example, it is lactaid pills.⁴

EffectiveExposure: The level of exposure to a hazard taking into account any mitigations. In this case, the lactaid increases Michelle’s susceptibility level to one scoop of ice cream.

Consequence: The effect of an object’s exposure to a hazard. Similar to Realization in [5]. Here, a stomach ache.

The remaining classes, **HazardousSituation**, **HazardousEvent**, and **Cause** remain as described in [1].

4.2 Axioms

We maintain all of the axioms in the current HazardousSituation ODP except that

$$HazardousEvent(he) \rightarrow \exists exposure(e) \wedge hasQuality(he, e)$$

⁴ <https://www.drugs.com/cdi/lactaid.html>

is removed and

$HazardousEvent(he) \rightarrow \exists hazard(h) \wedge hasParticipant(he, o) \wedge isExposedTo(o, h)$
becomes

$HazardousEvent(he) \rightarrow \exists exposure(e) \wedge hasParticipant(he, o) \wedge isExposedTo(o, e)$
to account for the change described in the second paragraph of this section.

In addition, we add the following additional axioms:

- A mitigation does not change the hazard an exposure involves.
 $mitigatesFrom(m, e1) \wedge mitigatesTo(m, e2) \wedge ofHazard(e1, h) \rightarrow ofHazard(e2, h)$
- Domain and range restrictions for all properties (as shown in Figure 1)
- Disjoint relations exist between all pairs of classes except Cause, Consequence, and HazardousEvent (e.g. one hazardous event can cause another).
- In situations in which a consequence can only be caused by exposure to one hazard, we can add an axiom expressing that if an object experiences a consequence from an exposure to a hazard, it must have been susceptible at that level of exposure. However, this is cannot be represented in OWL due to the existence of existentials on the right hand side of the rule.
 $experiencesConsequence(o, c) \wedge isExposedTo(o, e1) \wedge ofHazard(e1, h) \wedge hasAmount(e1, a1) \rightarrow$
 $\exists Susceptibility(s) \wedge hasSusceptibility(o, s) \wedge \exists Exposure(e2) \wedge toExposure(s, e2) \wedge$
 $\exists Amount(a2) \wedge ofHazard(e2, h) \wedge hasAmount(e2, a2) \wedge exceeds(a1, a2)$

The pattern is available on the Ontology Design Patterns website at <http://ontologydesignpatterns.org/wiki/Submissions:ModifiedHazardousSituation>. In addition, pattern development is hosted at Github⁵ in the Vocamp organization repository. Community comments and contributions are always welcome.

5 Conclusions and Future Work

In this work we stress the importance of modeling the consequences of exposure to hazards in a way that permits risk assessment and mitigation planning. Because current ontologies related to hazard exposure take an implicitly retrospective approach, we have proposed modifications to one such ontology, the Hazardous Situation ODP, to support these types of analysis. Our modifications draw inspiration from ontological models of biological dispositions.

There are some limitations to this pattern that we hope to address in our future work on this topic. One of these is that the modified version of the hazardous situation pattern is not completely compatible with the original, due to changing the range of the *involves* property from *Hazard* to *Exposure*. This change could arguably be made to the original pattern as well. The creators of that pattern define a hazardous event as “an event where at least one participating Object is exposed to a Hazard” [1]. The mention of an exposure in this definition seems to imply that a hazardous event actually involves an exposure to a hazard, not just the presence of one. Another issue is that, while the original version of the pattern supports the representations of actual hazardous situations in an a posteriori sense and the modified version of the pattern presented here supports a

⁵ <https://github.com/Vocamp/Hazard>

priori analysis, it is not currently clear how to connect the two. In other words, it is not clear how to connect information about an object’s susceptibility to a potential exposure to a hazard to information about an actual hazardous event that involved that exposure, such that meaningful inferences can be drawn. Further work needs to be done to enable this capability in a computationally feasible manner.

We plan to utilize this pattern in two very different applications: incident command response and sustainable building construction. In the incident command response project, the pattern will be used to represent hazards posed to people from chemicals, protective measures taken to protect individuals, and injuries sustained. The pattern will be populated with information from materials safety data sheets and emergency medical services “run sheets,” which describe a person’s injuries when an ambulance is called. For the sustainable-resilience application, the pattern will be applied to databases of building components to understand the frequency of replacement of building components under different hazard scenarios. For example, FEMA’s “P-58 Seismic Performance Assessment of Buildings”⁶ program defines fragilities, replacement probabilities and costs for different materials corresponding to different levels of seismic hazard.

Acknowledgments. The authors would like to thank the anonymous reviewers for their excellent insights regarding this pattern. They also gratefully acknowledge contributions from Civil Engineering domain experts Tracy Kijewski-Correa and Alexandros Taflanidis, the Center for Research Computing (CRC) at the University of Notre Dame and also the Green Scale Project. The authors would also like to sincerely thank the members of the Intelligent Agent ICS Augmentation team: Subhashini Ganapathy, John Gallagher, Jim Gruenberg, and Jack Smith. Vardeman and Ferguson would like to acknowledge support under National Science Foundation award CMMI 15-37652 “Life-cycle Assessment of Resiliency and Sustainability of Buildings.” Cheatham acknowledges support under National Science Foundation Award Number 1528550 “CPS EAGER: Intelligent Agent Incident Command System Augmentation.”

References

1. Lawrynowicz, A., Lawniczak, I.: The hazardous situation ontology design pattern. In: Proceedings of the 6th Workshop on Ontology and Semantic Web Patterns (WOP 2015). pp. 43–53 (2015)
2. Letia, I.A., Groza, A.: Developing hazard ontology for supporting haccp systems in food supply chains. In: IEEE 8th International Symposium on Intelligent Systems and Informatics. pp. 57–62. IEEE (2010)
3. Maalel, A., Mejri, L., Ghezala, H.H.B., et al.: Development of an ontology to assist the modeling of accident scenarii: Application on railroad transport. arXiv preprint arXiv:1203.1021 (2012)
4. Ortmann, J., Kuhn, W.: Affordances as qualities. In: FOIS. pp. 117–130 (2010)
5. Röhl, J., Jansen, L.: Representing dispositions. Journal of Biomedical Semantics 2(4), 1 (2011)

⁶ <http://www.fema.gov/media-library/assets/documents/90380>

6. Tierney, K., Lindell, M., Perry, R.: Facing hazards and disasters: Understanding human dimensions (2006)
7. Winther, R., March, W.: Hazards, accidents and events - a land of confusing terms (2013)
8. Wu, C.g., Xu, X., Zhang, B.k., Na, Y.l.: Domain ontology for scenario-based hazard evaluation. *Safety science* 60, 21–34 (2013)