Towards an Ontological Framework for Environmental Survey Hazard Analysis of Autonomous Systems

Dr Christopher Harper, Prof Praminda Caleb-Solly[1]*
[1]Bristol Robotics Laboratory, University of the West of England
T-Block, Frenchay Campus, UWE Bristol,
Bristol BS16 1QY, United Kingdom
chris.harper@brl.ac.uk, praminda.caleb-solly@uwe.ac.uk

Abstract

This paper presents current progress in the development of Environmental Survey Hazard Analysis (ESHA), a method of preliminary hazard identification aimed at autonomous system application problems. In addition to performing their design mission, autonomous systems must be capable of reliable and predictable behaviour in their environments, particularly when facing potential hazards that are not explicitly included in their design specifications ('non-mission' tasks). ESHA differs from conventional hazard identification methods in that its scope explicitly covers the identification of non-mission interactions between a system and its environment and any associated hazards. Although of general use as a safety analysis technique, ESHA has been designed primarily to support a "so far as is reasonably practicable" (SFAIRP) style of safety argument. However, early versions of the method were based on informal models, and therefore provided only weak support. This paper reviews the development of a formal ontological framework for ESHA, intended to provide much stronger basis for arguing the completeness and consistency of analyses.

1 Introduction

Autonomous systems (AS) are now emerging from the research laboratory into full industrial and social application, for example applications which require collaborative interaction with humans in shared spaces. Yet one of the key challenges in the development of these systems, and one of the principal barriers to their full deployment, remains unsolved; we still lack the tools and methods for adequate safety assurance and certification. This paper reviews current progress at Bristol Robotics Laboratory in the development of a method called Environmental Survey Hazard Analysis (ESHA) (Harper et al. 2014), which is a relatively new technique of preliminary hazard identification aimed at autonomous system design problems.

*Supported by the Assistive Robotics in Healthcare project of the Assuring Autonomy International Programme (www.york.ac.uk/assuring-autonomy/). We thank our colleagues Daniel Delgado Bellamy, Sanja Dogramadzi, Alex Sleat and Jason Welsby for their support.

1.1 The Problem of Autonomy from a Safety Perspective

Autonomous systems have unique characteristics and requirements that present a considerable challenge for safety assurance and certification. By definition, an AS may be required to operate for extended periods (or even indefinitely) without any human intervention or supervision. This means that they must be capable of interacting safely with any feature of the environment necessary to ensure their ongoing survival, and performance of their required mission. Hence, there is a set of "non-mission" interactions that an AS is required to perform, which relate to general existence and survival, as well as all those "mission tasks" that are required for the AS to fulfil its intended purpose.

A simple example of the concept of mission and non-mission interactions of an autonomous robot waiter is shown in Figure 1.

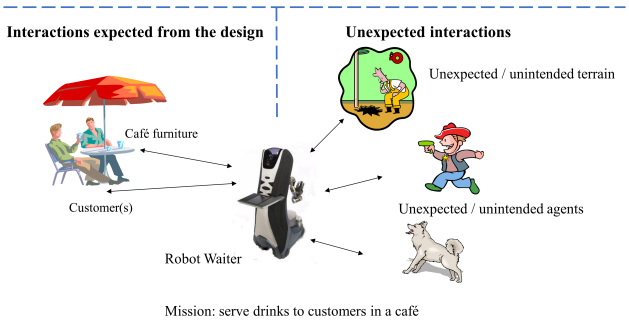

Figure 1: Mission vs. Non-mission Interactions

The robot will be expected to (and designed to) perform tasks related to interacting with customers, or setting down or picking up drinks from furniture such as tables. However,

there are numerous other unexpected features of the environment that might be found within a location where such a robot might operate(e.g. a cafe), for example unusual terrain (such as holes in the floor), humans such as children who are not otherwise engaged in the business of ordering drinks, or other creatures such as pets or service animals (e.g. guide dogs for visually impaired people). The robot will need to be capable of performing these non-mission interactions safely and reliably if it is not to present an unacceptable risk while in service.

Non-mission interactions are often overlooked by typical design engineering practices. However, this gap is usually closed either by the supervision of human operators or by constraining the system's environment, which reduces the number of required interactions between a system and its environment to a point where it becomes tractable to control them with the resources available to the system. Since such interactions are an essential feature of AS problems, we need new or revised safety analysis methods that seek to identify non-mission interactions, especially in unbounded environments. Such a goal is challenging - the unboundedness criterion creates problems of combinatorial expansion in the number of situated states that must be considered. We encountered these problems when we first started to look at hazard identification of AS, which led to the development of our original version of ESHA, as presented in (Harper et al. 2014).

1.2 Safety Validation of Autonomous Systems and ML/AI

In safety critical systems, hazard analysis is one of the principal sources of information for safety validation requirements, and autonomous systems are no exception in this regard. Once a hazard has been identified, it follows that one or more safety measures need to be introduced (typically design changes, safeguard mechanisms, or operational procedures) to reduce or eliminate the risk of its occurrence, and the effectiveness of those safety measures must be validated. Since ESHA is a process of identifying environmental interactions, it lends itself naturally to the specification of test scenarios that can serve to evaluate the effectiveness of safety measures. For example, it could be used in conjunction with scenario definition languages such as M-SDL (Foretellix 2020), OpenSCENARIO (ASAM 2020) or SCENIC (Fremont et al. 2019), where ESHA captures the high-level specification for scenarios that can be translated into scenario description languages for execution on a simulator.

Scenario-based validation is seen as one of the principal modes of safety validation for autonomous systems (see (Fremont et al. 2019) for example). As discussed earlier, one of the distinguishing concepts of autonomy relates to the requirement to interact reliably with features of the environment. So, it follows that one must test those interactions in order to validate the correctness of an autonomous system's function. Since ESHA is a process of systematic search for (potentially hazardous) environmental interactions, irrespective of their inclusion in the system's specified mission, it is a straightforward extension of the method to identify one or more validation scenarios for each identified hazard, to evaluate the capability of the AS to avoid them. ESHA is a complementary technique to a test coverage metric called *Situation Coverage* which has been developed recently (Alexander, Hawkins, and Rae 2015) as a validation metric for scenario-based testing, and in Section 5.2 we discuss current activities using ESHA to identify test scenarios for a connected autonomous vehicle (CAV) using situation coverage to evaluate progress and completion of the process.

The environment-driven requirements analysis produced by ESHA can also support the development of machine learning solutions for autonomous system problems. Since all ML processes or devices are developed inductively by training from a set of examples, the correctness of the system behaviour produced from a machine learning process is highly dependent on the diversity of the samples provided for its training. If there is any systematic error of omission in the training set, then the product of the ML process is likely to exhibit similar omissions in its operational behaviour. By providing a systematic coverage of potentially hazardous environmental interactions, ESHA can assist in specifying the scope of the training data set that is needed for adequate coverage of potential hazards, and can support the safety argument that an ML/AI system has been sufficiently well trained that the risk of any incompleteness affecting safety has been reduced as far as reasonably practicable.

1.3 Safety arguments for autonomous systems

The unboundedness of the general problem of autonomy (Pfeifer and Scheier 1999) causes significant problems for the safety argument(s) of an AS - for example as recommended by the new UL 4600 standard for autonomous road vehicles (ANSI/UL 2020). The demonstration that the residual risk of harmful events (a function of severity and probability) of a system is 'acceptable' becomes questionable. Significant underlying causes of this problem include:

1. In unbounded environments, the number of features with which an autonomous agent may interact is uncertain. This affects the validity of quantitative risk analyses; calculations of requirements for or verification of a system's probability of failure depend on the set of contributing hazards to be complete. If any are omitted then quantitative safety targets will be erroneous and the system less reliable or fail-safe than it needs to be to fulfil its 'true' requirements (i.e. those specified by legislation).
2. It has been a long standing problem in the risk assessment of safety critical systems in many industry sectors, that regulatory requirements demand levels of reliability that are not feasible to demonstrate by practical testing (Butler and Finelli 1991). Most practical test programmes can only deliver high-confidence estimates of probability of failure that are usually several orders of magnitude lower than the values required. This applies in particular to rare events, which are unlikely to occur during testing but may nevertheless still occur at a higher rate than is considered acceptable by legislative requirements.
3. Where simulation models are used to accelerate the testing of a system, and existing simulation data are used for

probability estimation based on extrapolation of (bounds of) probability from the measured results, any such analysis rests on the assumption that the simulation is a faithful reproduction of the system and its environment. But any simulator will have limits to its fidelity (Koopman and Wagner 2018), which are likely to be impossible or too expensive to correct, and this makes the task of identifying rare events solely from testing impossible (the event does not occur during the test programme, and there is no way to determine whether the lack of occurrence is due to the genuine rarity of the event or due to being masked by errors in simulator fidelity).

In all these cases a systematic analysis of environmental interactions such as ESHA will help to improve confidence in the quality of the results, but our conclusion is that it will be difficult to produce strong supporting evidence for an AS safety case with these techniques.

Therefore, we are investigating an alternative safety argument strategy (Harper 2020), where the claim is that the risk has been reduced "so far as is reasonably practicable" (SFAIRP). This argument concept is explicitly written into UK law (HMSO 1974) and is often an acceptable legal interpretation in other countries. We use the UK legal basis of the concept as the working model for our studies, as well as the work in (Eliot 2006, 2007), which provides supporting interpretive material.

In a SFAIRP-based safety case, the general obligation for a system developer is to demonstrate that any of the following have been achieved:

- it was not practicable or not reasonably practicable to do more than was in fact done to satisfy the duty or requirement;
- there was no better practicable means than was in fact used to satisfy the duty or requirement;
- the potential cause of harm lies *outside the scope of the [system developer's] undertaking* (Eliot 2007).

Compliance with these requirements must be demonstrated by means of *objective* evidence, namely that no allowance be made for personal qualities of any legal defendant, i.e. it will not be a valid legal defence if the system developer appeals to the actual conduct of the safety assurance process, or the nature of the people who performed it. Evidence must be objective, i.e. epistemically independent of the people who created it.

The legal concept of "reasonableness" is based on the concept of a "reasonable person", which in the case of engineering problems may be interpreted as "competent technical specialist". Hence whatever may be considered 'reasonably practicable' in a safety assurance context would mean whatever might be expected of a suitably qualified practicing safety assurance engineer. If in the design of any system, any hazard was overlooked or any safety feature of a system design not considered for its practicability, which could plausibly have been identified or specified by such an engineer, then it could not be said that all foreseeable hazards were identified, nor all practicable safety measures deployed.

A plausible defence against breach of safety regulations can be made if it can be demonstrated *objectively* that the harm was not foreseeable, any safety measures not deployed were not reasonably practicable, or the harm was outside the scope of the undertaking. For these reasons, a safety case/argument will need to demonstrate objectively (Eliot 2007) the following characteristics:

- that exhaustive coverage of potential hazards has been achieved
- that the analysis of a system's scope of operation is systematic and complete
- that the identification of safety measures to resolve any potential hazard has been exhaustive
- that the costs vs. benefits of each identified safety measure have been analysed systematically

However, this may not be easy for AS - if the environment is unconstrained or unbounded, then it may not be possible to show that any potential cause of harm is outside the scope of the undertaking, or that it was not practicable to do more than was in fact done to identify potential hazards.

ESHA was developed (Harper et al. 2014) specifically to support this type of argument by attempting to construct at least part of the evidence necessary to support such claims, in terms of providing an objectively demonstrable process framework that can assist analysts to show that all *categories* of environmental feature have been considered, even if it is not feasible to demonstrate that all *instances* have been identified. This is the motivation for our consideration of ontological frameworks; ontology is the disciplined (and, ideally, formal) practice of trying to build such classification schemes.

While the SFAIRP approach takes a different approach by avoiding conventional quantitative risk-driven arguments, it is nevertheless the case that reducing risk to as low as can practicably be achieved may still mean that the residual risk is too great to be acceptable. This is a residual problem, but not one that can be resolved by engineering analysis methods alone; it is a legislative matter that lies within the purview of legislators and politicians to decide where any such boundary may lie. Nevertheless, we argue that ESHA could contribute to such decision making by supporting the establishment of risk models that have some formal backing, and some confidence that the significant risks have been found.

2 Overview of ESHA

ESHA was developed to help identify environmental features with which an AS might interact (irrespective of whether it is a mission or non-mission interaction) and whether there is any potential for hazard.

2.1 ESHA Procedure and Guide-words

The ESHA procedure (Harper et al. 2014) is similar in many respects to conventional system hazard analyses: a set of *guide-words* is considered, which classify the environmental features with which an AS may need to interact. The general nature of any possible hazardous interactions is identified. The results are compiled in tabular a format similar to traditional variants of hazard analysis. A set of procedures and

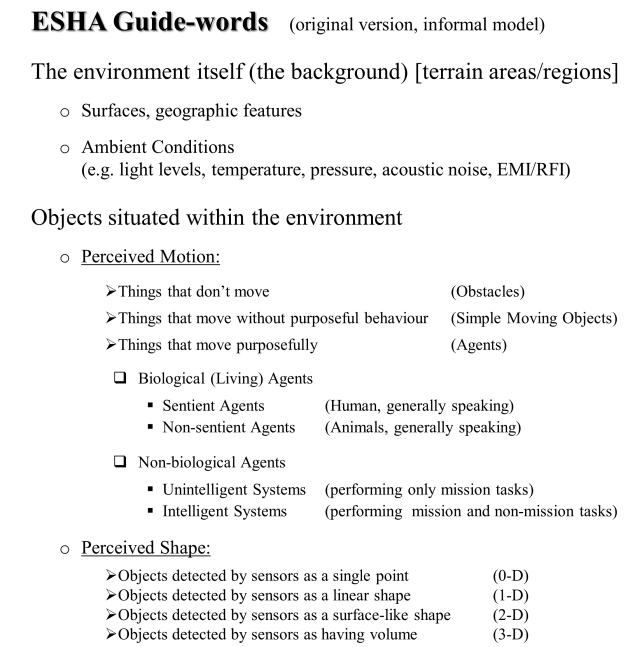

Figure 2: Original ESHA Guide-words

checklists were developed to guide analysts in compilation of results tables correctly, and the guide-words shown 2 in were specified.

While this guide-word set is not unreasonable in an intuitive sense, and was intended to be grounded in a concept based on how the characteristics of how environmental features might be perceived by the sensors of an AS, it nevertheless is an informal model, and lacks any objective proof of completeness. This compromises the aim that the ESHA method can demonstrate objectively an exhaustive coverage of interactions with environmental features. We are therefore looking to provide such a formal basis, which is the reason for our investigation of ontological frameworks as a basis for doing so.

2.2 Current Experience with ESHA

Our applications of ESHA have so far been restricted to small or partial application problems in robotics and autonomous systems. In the original development phase of ESHA (Harper et al. 2014), we investigated some problems in urban search and rescue and domestic assistance (guide robot for elderly persons). More recently, we have trialled the method in conference and project workshops, including the SOCRATES project[1] and the European Robotics Forum 2020 (Caleb-Solly 2020).

Anecdotal and verbal responses from workshop delegates and participants have generally been positive, although it should be noted that in most cases the participants involved were not safety assurance practitioners (academic researchers and students in robotics were the most typical groups). And even though these versions of ESHA have been logically informal, the method still serves as a useful technique for getting human designers to consider in a structured way how hazards might be identified for autonomous systems. But it has been noted that complexity and combinatorial issues may limit the practicability of the method, and ontology has been seen as one of the best approaches to overcoming these problems.

[1]This event is mentioned in passing at URL: http://www.socrates-project.eu/blog/2019/12/02/meeting-in-bristol/

3 Requirements for an ESHA Ontology

The guide-words presented in the previous section are an attempt to categorize environmental features in such a way as to be logically complete and therefore exhaustive of any environmental domain. The set was developed on the conceptual basis that in an unbounded environment, the only boundary available from which one could extract a logically complete description was the boundary of the system (robot) itself. The intent of the original ESHA guide-words was to identify environmental features by the characteristics of their 'images' as they impinged upon the system sensors. Visual perception was the principal modality considered, hence the classification of objects in particular as 'point-like' or 'line-like'; the 'dimensionality' of images was considered to be a concept that could demarcate all possible environmental features in a logically exhaustive way.

However, the guide-word model was never formally derived from first principles, and hence there is no rigorous basis for asserting its completeness. Since we wish to demonstrate objectively that the ESHA procedure is exhaustive, we are investigating existing formal ontologies to see if they provide the necessary logical framework we need to underpin ESHA. Based on the previous discussions, the following requirements are seen as necessary for the foundational ontology:

- Support for essential ontological concepts
  Several core ontological topics are required a a minimum of any ontology that could be used to support ESHA:
  - *Mereotopology*: relationships between parts and wholes, that allow the identification of structures and complex objects
  - *Boundaries*: formal treatment of boundaries within the mereotopology, especially *fiat* as well as *bona fide* boundaries (Smith and Varzi 2000)
  - *Situations and Events*: formal classification scheme of situations and events, which will provide support for hazard identification
  - *Agency, Causality, Autonomy*: formal classification scheme for the behaviour of environmental features, to allow capture of causal relationships and identification of interactions.
- Particulars *and* Universals
  Since we are interested in establishing abstract concepts such as safety properties, we seek a foundational ontology that incorporates Universals as well as Particulars.

- Realism or Conceptualism preferred over Nominalism
  Since we require the existence of universals, we reject the purely Nominalist stance that Universals do not exist. We require ontological frameworks that reflect at least a Conceptualist if not a fully Platonic perspective. However, this does not mean that we cannot incorporate nominalist ontological models as partial frameworks, applying solely to Particulars, and then 'complete the model' by adding corresponding Universals, which can be done by means of model-patterns such as the *ontological square*.
- Logical completeness and disjointness
  We are looking for formal ontologies whose type hierarchy is defined as far as possible in a logically complete and disjoint manner. Completeness ensures that the model is exhaustive in its scope. Disjointness of types ensures that our analysis remains tractable, since multiple combinations of parent types will be excluded.

Our general model of ontological frameworks for ESHA takes the view that they will follow a three-layer organization/ structure:

- *Foundational ontologies* are the most abstract layer, defining the most basic entities that underpin other layers.
- *Domain ontologies* define the most general concepts that are specific to a particular domain but general to numerous applications (e.g. system safety, geospatial domains, HMI, etc.);
- *Application ontologies* define entity types intended for specific application categories, such as assistive robots or driverless vehicles

4 Review of Candidate Ontologies

It was not the original aim of this work to develop any fundamentally new ontologies to support ESHA, rather to exploit existing work (Harper 2020). However, in respect of the basic *foundational ontology* layer, no existing frameworks have been found to possess all the properties desired. So we must adapt existing foundational ontology to introduce the modifications or new elements as needed.

While numerous proposed ontological frameworks have been reviewed, the following are the major candidates that were considered for adoption as the ESHA foundational ontology.

- Sowa's Knowledge Representation Ontology (Sowa 2000)
- Basic Formal Ontology (Smith 2015)
- Zemach's "Four Ontologies" (Zemach 1970)
- Unified Foundational Ontology (UFO) (Guizzardi 2005)

Sowa's ontology (Sowa 2000) was the first we reviewed. It blends together several ontological concepts developed by Peirce and Whitehead, and has some interesting features, but two significant flaws. First, as noted by (Degen et al. 2001), it does not draw clear distinctions between sets, universals and individuals, nor does it clarify the ontological meaning of modal operators used in their definitions. Second, we have noticed that the type hierarchy contains inheritance errors in some sub-types, which are derived both from physical and abstract entities simultaneously and therefore inherit conflicting definitions of spatial and temporal existence. Once the type hierarchy is corrected for this inconsistency, it begins to look similar to other ontologies (such as BFO, discussed below) that are far more explicitly formal and hence are preferred for that reason.

Basic formal ontology (Smith 2015) has the advantage of being a consistent formal specification, but is insufficient for ESHA purposes since (as a matter of pragmatic policy) it only include types and sub-types of Particulars. Additionally, the theory of mereology that underpins the BFO model is less extensive than other frameworks (UFO).

Zemach's "Four Ontologies" model (Zemach 1970), although having a nominalist perspective, does offer a disjoint and complete decomposition of Particulars into four types of continuant and occurrent entity. The model is complete and disjoint, and we believe can be adapted to resolve issues in other models to improve the properties of the model eventually used to support ESHA (as discussed later).

While no previously developed formal ontological model had all the attributes we were seeking, our preferred ontological framework is the Unified Foundational Ontology (UFO), as it has the greatest number of useful features that we have seen, and we believe that the few deficiencies of logical completeness that exist in some parts of its model (especially in the decomposition into continuants and occurrents) can be modified to correct the problems. While the original version of UFO (UFO-A) was developed only as an ontology of endurant (continuant) entities (Guizzardi 2005), it has been extended with additional subsets UFO-B (covering perdurant/occurrent entities) and UFO-C (providing models of causality, interaction and agency). UFO-A and UFO-B have recently been integrated into a combined ontology of endurants and perdurants UFO-AB (Benevides, Almeida, and Guizzardi 2019). These extensions develop ontologies for processes, events and their relationships, for example temporal sequences of events, or how processes might be composed from (or otherwise related to) set(s) of individual events that may be ascribed to them.

5 Conclusion and Current Progress

We have selected UFO as the basis for a design of a foundational ontology for ESHA, although it will require some modification and extension to ensure that it has the relevant properties of completeness necessary to fulfil the requirements of a method that can support a SFAIRP safety argument.

5.1 Modifying UFO to Support ESHA Requirements

Where the UFO ontological model in its most recent incarnation UFO-AB (Benevides, Almeida, and Guizzardi 2019) has all the properties we require for ESHA, we propose to use it unchanged. Where it does not, we propose to incorporate ideas from other models where they appear to be compatible, a partial example of which is shown in the model fragment in Figure 3.

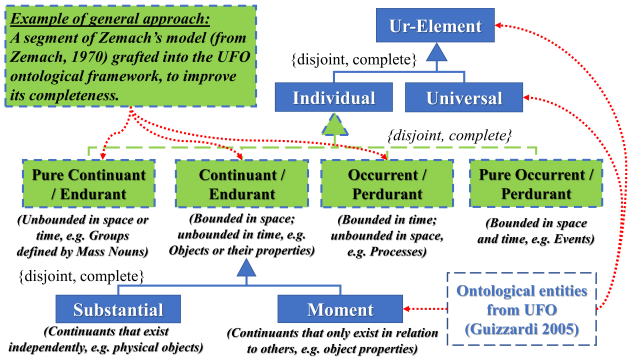

Figure 3: Enhancement of UFO model (where necessary) by incorporation of elements from other frameworks

Figure 3 shows a partial representation of one particular modification that we have already identified, using Zemach's model mentioned above. In the original type model developed for UFO, the major sub-types of Individual consisted only of 'Endurant' (also known as 'Continuant') and 'Perdurant' (also known as 'Occurrent'), and this was not a logically complete and disjoint type decomposition. By adding two new sub-types 'Pure Continuant' and 'Pure Occurrent' (known as 'Event' in Zemach's model) one can develop this into a complete and disjoint decomposition, albeit with some underlying logic regarding spatial and temporal bounds, as mentioned in the list below.

By this and other modifications, we aim to create a new *ESHA Foundational Ontology* as an evolutionary development of UFO. There are several 'grafts' to be done to UFO to incorporate all the elements that we anticipate might be necessary or useful to support an environmental survey hazard analysis, including:

- The incorporation of Zemach's model as the major subtypes of Individuals, and the extension of this model (which was originally conceived as a nominalist model applicable only to Individuals, also known as 'Particulars') into the corresponding Universal types according to the ontological square, which UFO applies as a metamodel governing the form of its ontology.
- The incorporation of the concept(s) of *topoids*, *chronoids*, and *situoids* from GOL (Degen et al. 2001) as the underlying formal framework for the elements from Zemach's model.
- The incorporation of *Fiat Boundaries* (Smith and Varzi 2000) into the mereotopology framework of UFO; *fiat* boundaries are boundaries that are established by an observer or social convention rather than existing in a strictly physical sense. This concept may be useful because changes in agent behaviour can be driven by such boundaries (for example a doorway between two rooms, or the centre line of a road), and interactions between agents and the environment may be influenced or even governed by such boundaries.
- The incorporation of the work of (Vogt et al. 2011), who have developed a complete taxonomy of constitutively organized material entities as an extension of BFO. This work could be a bridge between the foundational ontology and domain ontology layers, as it could form an abstract catalogue of patterns for combining foundational material entities in any given domain, which can in principle be logically complete.
- It may be worthwhile to incorporate some of the work of Bittner on granular partitions, collections, and temporal mereological relations (Bittner, Donnelly, and Smith 2006). Bittner's work on partitions may also serve as a meta-model for the development of the ontology itself, as it can serve to establish the consistency and completeness of the type hierarchy.

The extent to which it is logically permissible to incorporate all these elements in a consistent manner remains to be seen, and will be a challenge for this research programme. Where it proves impossible to complete the UFO ontology, 'gaps' will remain in the underlying framework, which then affect the capability of environmental survey analyses to conclude that a systematic and complete survey has been achieved. We anticipate (purely as an informal judgement) that it will be possible to develop a complete Foundational Ontology for ESHA, but that incompleteness may begin to 'creep in' to the model from the Domain Ontology level onward (as discussed in Section 3), where the relatively abstract types of the foundational layer begin to be applied to real-world domains.

For example, the formal ontological 'catalogue' of constitutive material entities developed by (Vogt et al. 2011) would seem to be an extremely powerful development of profound significance to this research, but early indications are that for any given domain the catalogue may be extremely large due to the permutations of potential combinations of entities, and may well be open-ended. (Consider the set of all possible roads that can be composed from road environment entites such as junctions, straight sections, road bends, etc.) Hence the domain-level catalogue may become impracticable to complete, even though a complete set of individual entities might be identified.

But the advantage of attempting to ground the analysis in an ontological model means that at least it will be known where the incompleteness exists, and this can be taken into account when identifying hazards and specifying safety requirements for a system. (Effectively, this is a strategy of attempting to transform any "unknown unknowns" of a hazard analysis into "known unknowns", which can then at least be managed if not resolved.)

5.2 Current Experiments and Further Work

We are currently applying the improved ESHA methodology to application problems in the fields of assistive robotics for healthcare and social care, and also to driver-less road vehicles (CAVs).

In the CAV domain, we are using the ESHA technique to analyse the operating design domain (ODD) of a connected autonomous vehicle (CAV) in order to develop a specification of test scenarios based on a systematic review of the vehicle route. We will generate CAV simulator test scenarios

and validate coverage of the domain by use of a validation metric called *situation coverage*, which is a new validation metric concept developed recently by (Alexander, Hawkins, and Rae 2015).

Our experiments in assistive robotics applications are aimed at gaining broader experience in application of the method, especially since many applications in this field are concerned with human-robot interaction with vulnerable users with complex care needs that vary over time. This increases the complexity of the required interactive behaviour both for mission and non-mission tasks, and we are investigating the impact of this on the practicability of the method.

We will then proceed to develop domain-level and application-level ontologies that support two particular domains (assistive medical/social care robots, and driverless road vehicles), which are active research topics at Bristol Robotics Lab. As each stage of the ontological framework is completed it will be published as a journal paper, and we plan in the long term to amalgamate all the work into an ESHA Handbook, to be made publicly available as a textbook.

References

Alexander, R.; Hawkins, R.; and Rae, A. J. 2015. Situation coverage – a coverage criterion for testing autonomous robots. Technical Report Technical Report YCS-2015-496, Department of Computer Science, University of York.

ANSI/UL. 2020. *Standard for Safety for the Evaluation of Autonomous Products.*, ansi/ul standard 4600, 1st edition.

ASAM. 2020. *OpenSCENARIO Manual*, v1.0.0 edition. URL https://www.asam.net/standards/detail/openscenario/.

Benevides, A. B.; Almeida, J. P. A.; and Guizzardi, G. 2019. Towards a Unified Theory of Endurants and Perdurants: UFO-AB. In *Proceedings FOUST III: Workshop on Foundational Ontology, The Joint Ontology Workshops (JOWO 2019)*, volume CEUR Workshop Proceedings Vol. 2518. Graz, Austria. URL http://ceur-ws.org/Vol-2518/.

Bittner, T.; Donnelly, M.; and Smith, B. 2006. A Spatio-Temporal Ontology for Geographic Information Integration. *International Journal of Geographical Information Science* 23(6): 1–29.

Butler, R. W.; and Finelli, G. B. 1991. The Infeasibility of Experimental Quantification of Life-Critical Software Reliability. In *ACM Software Engineering Notes (Proc. SIGSOFT '91 Conf. on Software for Critical Systems)*, volume 16(5), 66–76. New Orleans.

Caleb-Solly, P. 2020. ERF 2020 – Workshop Report: Assuring Safety for Assistive Robotics in Health and Social Care. URL https://www.eu-robotics.net/robotics_forum/upload/erf2020/presentations/Workshops_04.03.2020.rar.

Degen, W.; Heller, B.; Herre, H.; and Smith, B. 2001. GOL: Towards an Axiomatized Upper-Level Ontology. In *Proceedings FOIS'01*, volume ACM 1-58113-377-4/01/0010. Ogunquit, Maine, USA.

Eliot, C. 2006. System safety and the law. In *Proceedings 1st IET International Conference on System Safety*, 344–351. London, UK. ISBN 0-86341-646-2.

Eliot, C. 2007. What is a reasonable argument in law? In *Proceedings of 8th GSN User Club Meeting*, 344–351. York, UK. ISBN 0-86341-646-2.

Foretellix. 2020. *Open Measureable Scenario Description Language Manual*, v20.10 edition. URL https://www.foretellix.com/open-language/.

Fremont, D. J.; Dreossi, T.; Ghosh, S.; Yue, X.; Sangiovanni-Vincentelli, A. L.; and Seshia, S. A. 2019. Scenic: A Language for Scenario Specification and Scene Generation. In *Proceedings of the 40th ACM SIGPLAN Conference on PLDI '19*. Phoenix, AZ, USA. doi:https://doi.org/10.1145/3314221.3314633.

Guizzardi, G. 2005. *Ontological Foundations For Structural Conceptual Models*. Ph.D. thesis, Centre for Telematics and Information Technology, University of Twente, The Netherlands.

Harper, C. 2020. Environmental Survey Hazard Analysis: Current Developments. UK Safety Critical Systems Club seminar, *New Safety Analysis Techniques*. URL https://scsc.uk/e654.

Harper, C.; Dogramadzi, S.; Giannaccini, M. E.; Sobhani, M.; Woodman, R.; and Choung, J. 2014. Environmental hazard analysis - a variant of preliminary hazard analysis for autonomous mobile robots. *Journal of Intelligent and Robotic Systems* 76(1): 73–117.

HMSO. 1974. *The Health and Safety at Work etc Act*. URL https://www.legislation.gov.uk/ukpga/1974/37/contents/.

Koopman, P.; and Wagner, M. 2018. Toward a Framework for Highly Automated Vehicle Safety Validation. In *Proceedings of SAE World Congress 2018*, 1–13.

Pfeifer, R.; and Scheier, C. 1999. *Understanding Intelligence*. MIT Press. ISBN 0-262-16181-8.

Smith, B. 2015. *Basic Formal Ontology 2.0: Specification and User's Guide*. URL https://basic-formal-ontology.org/.

Smith, B.; and Varzi, A. C. 2000. *Fiat* and *Bona fide* boundaries. *Philosophy and Phenomenological Research* LX(2): 401–420.

Sowa, J. F. 2000. *Knowledge Representation: Logical, Philosophical, and Computational Foundations*. Pacific Grove, CA: Brooks Cole. ISBN 0-534-94965-7.

Vogt, L.; Grobe, P.; Quast, B.; and Bartolomaeus, T. 2011. Top-Level Categories of Constitutively Organized Material Entities - Suggestions for a Formal Top-Level Ontology. *PLoS ONE* 6(4): 1–14.

Zemach, E. 1970. Four Ontologies. *Journal of Philosophy* 67(8): 231–247.