

Using ontologies to aid navigation planning in autonomous vehicles*

CRAIG SCHLENOFF,¹ STEPHEN BALAKIRSKY,¹
MIKE USCHOLD,² RON PROVINE² and SCOTT SMITH²

¹Nat. Inst. of Standards and Technology, 100 Bureau Drive, Stop 8230, Gaithersburg, MD 20899, USA; e-mail: craig.schlenoff@nist.gov

²Boeing, Phantom Works, P.O. Box 3707, m/s 7L-40, Seattle, WA 98124-2207, USA; e-mail: michael.f.uschold@boeing.com

Abstract

This paper explores the hypothesis that ontologies can be used to improve the capabilities and performance of on-board route planning for autonomous vehicles. We name a variety of general benefits that ontologies may provide, and list numerous specific ways that ontologies may be used in different components of our chosen infrastructure: the 4D/RCS system architecture developed at NIST. Our initial focus is on simple roadway driving scenarios where the controlled vehicle encounters objects in its path. Our approach is to develop an ontology of objects in the environment, in conjunction with rules for estimating the damage that would be incurred by collisions with the different objects in different situations. Automated reasoning is used to estimate collision damage; this information is fed to the route planner to help it decide whether to avoid the object. We describe our current experiments and plans for future work.

Introduction

An autonomous vehicle is an embodied intelligent system that can operate independently from human supervision. The field of autonomous vehicles is continuing to gain traction both with researchers and practitioners. Funding for research in this area has continued to grow over the past few years, and recent high-profile funding opportunities have started to push theoretical research efforts into practical use.

To behave appropriately in an uncertain environment, many researchers and practitioners believe that “the vehicle must have an internal representation (world model) of what it feels and experiences as it perceives entities, events, and situations in the world. It must have an internal model that captures the richness of what it knows and learns, and a mechanism for computing values and priorities that enables it to decide what it wishes to do” (Albus *et al.*, 2002, p. 196). The inability to do this well hinders effective task planning and execution and thus the overall effectiveness of the vehicle. A major challenge in autonomous vehicles is the ability to accurately maintain this internal representation of pertinent information about the environment in which the vehicle operates.

The field of autonomous vehicles has reached a level of maturity such that it could greatly benefit from leveraging the latest technologies in the area of knowledge representation and ontologies.¹ The

*Certain software tools are identified in this paper in order to explain our research. Such identification does not imply recommendation or endorsement by the National Institute of Standards and Technology, nor does it imply that the software tools identified are necessarily the best available for the purpose.

¹ The 2004 AAAI Spring Symposium series includes a workshop on this the topic: “Knowledge Representation and Ontology in Autonomous Systems”. See <http://www.aaai.org/Symposia/Spring/2004/sssparticipation-04.pdf>.

use of ontologies and automated inference is a natural fit for representing and reasoning about world models for autonomous vehicles. A large body of work exists in the ontology area, yet the authors are not aware of any that has been applied to the area of world-modelling in autonomous vehicles. This sets the context for our overall goal: to apply ontologies to improve the capabilities and performance of on-board route planning for autonomous vehicles.

The potential benefits of introducing an ontology (or set of ontologies) into an autonomous vehicle's knowledge base are many. One is the potential for reuse and modularity. For example, a general theory of obstacles could apply to a broad range of autonomous vehicles. In addition, ontologies provide a mechanism to allow for a more centralised approach for representing and reasoning about information about the environment. Different modules in an autonomous vehicle would query the ontology, rather than having the information scattered among the modules. This has a corresponding benefit in cheaper and more reliable maintenance. Finally, there is the potential for increased flexibility of response for the autonomous vehicle. Methods that rely on preclassification of certain kinds of terrain in terms of their traversability (Donlon & Forbus, 2003; Malyyankar, 1999) are important, but do not support reasoning about objects in a more dynamic context.

This paper proceeds as follows. We first discuss some related work. We then introduce the 4D/RCS autonomous system reference architecture that we are using to perform experiments. Section 4 takes a closer look at how and where ontology technology can be deployed in the various components of the 4D/RCS architecture. Subsequent sections present our early results on how to represent and reason about objects and obstacles to aid in vehicle navigation. We use some simple scenarios and an initial classification of collision damage severity to illustrate our ideas. Our goal is to create a general theory of obstacles that can be applied in a wide variety of situations. We finish with some ideas on future work and present our conclusions.

2 Related work

The most closely related work is in the area of determining "trafficability". This is defined to be "a measure of the capability for vehicular movement through some region" (Donlon & Forbus, 2003, p. 2), i.e. specific kinds of terrain. This work is being performed in the context of traditional GIS algorithms used for route planning, augmented with qualitative reasoning techniques (Forbus *et al.*, 2001). In the system described in Donlon & Forbus (2003), terrain is regarded as being in one of three categories: unrestricted, restricted or severely restricted. The idea is to pre-classify certain kinds of terrain in terms of traversability. Slope, hydrography and vegetation are also taken into account. For example, a path segment with a slope angle that is greater than 45 degrees would be severely restricted for most 4-wheel vehicles.

Similar work is reported by Malyyanker (1999), where the creation of a "navigation ontology" in a marine environment is discussed. It is also set in the context of GIS. We will study and mine these sources for ideas that we hope to generalise and apply. However, all of the reasoning used by these approaches relies on a priori information resulting from preclassifying known terrain. They do not address the issue of reasoning with *in situ* information received from the sensors during vehicle movement. Such dynamic classification capability will be a major focus of our research.

The work described in this paper also builds on previous work performed at the National Institute of Standards and Technology in the areas of applying ontologies for the purpose of object classification (Schlenoff, 2002) and some general thought papers on the role of ontologies in autonomous systems (Meystel *et al.*, 2002).

3 The 4D/RCS reference model architecture

We have selected the Real-Time Control System (4D/RCS) (Albus, 1991; Albus *et al.*, 2002) as the architecture in which we will implement and evaluate the use of ontologies for autonomous vehicles.




Figure 1 RCS Node

4D/RCS is a hierarchical, distributed, real-time control system architecture that provides clear interfaces and roles for a variety of functional elements.

Under 4D/RCS the functional elements of an intelligent system can be broadly considered to include behaviour generation (task decomposition and control), sensory processing (filtering, detection, recognition, grouping), world-modelling (store and retrieve knowledge and predict future states) and value judgement (compute cost, benefit, importance and uncertainty). These are supported by a knowledge database, and a communication system that interconnects the functional elements and the knowledge database. This collection of modules and their interconnections make up a generic node in the 4D/RCS reference model architecture (see Figure 1) (Albus & Meystel, 2001, p. 146). A generic node is defined as a part of the 4D/RCS system that processes sensory information, computes values, maintains a world model, generates predictions, formulates plans and executes tasks. Each module in the node may have an operator interface.

There are several contemporary architectures that exist in the literature for designing intelligent systems. We have several motivations for selecting 4D/RCS:

- In the last fifteen years, behaviourist architectures (Arkin, 1989; Brooks, 1985) have gained popularity for their ease of implementation. Reactive architectures tend to be unable to achieve long-range goals and to have additional problems with respect to scalability and reactions possibly overriding basic safety systems (Balakirsky, 2003).
- 4D/RCS is a proven architecture with many person-years of research and development in intelligent control theory. It has been implemented and tested thoroughly both in industry and academia in different operating domains under varying operating conditions. Initial applications were in civilian domains, over two decades ago. An application summary reference (Albus, 1995) describes 10 major systems from manufacturing machining stations to space robotics, automated coal mining, autonomous maneuvering control for submarines and underwater robots, and order fulfilment and stock materials handling control for stamp distribution for the US Postal Service. Examples of newer systems engineered based on RCS include a Postal Service large integrated mail tray handling system, automated process-adaptive riveting systems for wing assemblies at a major airplane manufacturer, a process-adaptive inspection and milling machining centre at a powertrain manufacturing plant and a mobile robotic manipulator for welding and painting of Navy double-hull ships. 4D/RCS has most recently been implemented as the reference model architecture for the design, engineering, integration and testing of eXperimental Unmanned Vehicles for the DoD Demo III programme (Albus, 2002; Gazi *et al.*, 2001).
- 4D/RCS is supported in terms of software and updates and thus it has continually evolved through a number of versions at National Institute of Standards and Technology (NIST) and elsewhere (Gazi *et al.*, 2001). For additional advantages, see pp. 128ff of Albus and Meystel (2001).

4 The role of ontologies

There are many challenges in the autonomous vehicle research community that ontology technology can help to address. In this section, we will look at the challenges faced by each of the functional elements shown in the previous section, and describe how adding an ontology-based component to the architecture adds value. The ontology component will contain the ontology and an associated inference engine. Its role is to provide answers to questions of various kinds that arise during vehicle navigation. In each subsection we state a set of competency questions (Gruninger & Fox, 1994) that the ontology component must answer to achieve the desired functionality.

4.1 Sensory processing

The sensory processing functions operate on data from sensors so that the world-modelling functions can maintain the knowledge database as a current and accurate estimate of the state of the world, including the internal state of the system itself.

Sensory processing is organised around sensors. Each sensor produces a signal that varies in time as the physical phenomena that it measures vary in time. Sensory processing functions operate on sensory signals to window, group, filter, compare, classify and interpret them as entities, events and situations that correspond, in a meaningful and useful way, to the actual entities, events and situations in the real world (Albus *et al.*, 2002).

The ontology component can play a substantial role in object classification and recognition by providing a sensory process system with the ability to reason over information that it discovers to determine which object class best matches the information that it currently knows. In this situation, the following competency questions arise:

- If I see an object with certain properties,
 - (a) What is it? What is it not?
 - (b) At what level of detail can I determine what it is? (e.g. is it a vehicle, a four-wheeled vehicle, a van, a minivan?)
 - (c) Is that level of detail enough to determine whether it is an obstacle, and to what extent?
 - (d) How confident am I that the object is what I determine it to be based on continuously updated sensor input and other reasoning?
- If I see an object with certain properties and I am not sure what it is, what additional information should I gather so that I will be better able to identify the object?
- If I see a group of objects that seem to form a particular situation (e.g. a “MEN WORKING” scenario), what additional objects should I be on the lookout for (e.g. men walking around)?

4.2 World-modelling

World modelling is a process that performs four principal functions:

1. It maintains a knowledge database of images, maps, objects, agents, situations, relationships and knowledge of task skills and laws of nature and relationships between them.
2. It maintains a best estimate of the state of the world to be used as the basis for predicting sensory feedback and planning future actions.
3. It predicts (possibly with several hypotheses) sensory observations based on the estimated state of the world. Predicted signals can be used by sensory processing to configure filters, masks, windows and templates for correlation, model matching and recursive estimation.
4. It simulates results of possible future plans based on the estimated state of the world and planned actions. Simulated results are passed to the value judgement system for evaluation in order to select the best plan for execution (Albus *et al.*, 2002).

Once an object in the environment is recognised from sensor data, the ontology can provide a wide variety of additional information about the object that is not directly perceivable (e.g. weight,

maximum speed, minimum turning diameter etc.). This information can be used to aid in determining the possible behaviours of the object and the potential damage from a collision. In turn, this information can be used in vehicle navigation planning. The following competency questions must be answered to provide adequate support in this situation:

- If I see an object of type X, then
 - (a) What is the range of possible speeds that it can be going?
 - (b) What are its possible directions of travel?
 - (c) What is the possible rate of change of direction of travel, at a given speed?
 - (d) Does it have weapons and am I in its range?
 - (e) Is it likely to be friend or foe?

The answer to these questions would often heavily depend on context. For example, if a man is standing holding a lollipop sign with slow or stop on either side, then he is more likely to move into traffic.

4.3 Value judgment/behaviour generation

The value judgement component evaluates perceived and planned situations. It computes what is important (for attention), and what is rewarding or punishing (for learning). The value judgement component assigns priorities and computes the level of resources to be allocated to tasks. It assigns values and costs to recognised objects and events, and computes confidence factors for observed, estimated and predicted attributes and states (Albus *et al.*, 2002). The outputs of the value judgement component are used by the behaviour generation component to select and set priorities during route planning.

We are exploring the use of ontologies as a mechanism to allow the planner in the behaviour generation component to understand better the costs and consequences of colliding with other objects. By representing the factors that could impact a path's cost, an ontology can be used to reason over the information that is available to determine what the consequences of a collision would be. Further reasoning could then be performed to determine the cost of these consequences (see Section 6 for more details). This cost would then be fed back to the planner for consideration when deciding the “cheapest” plan for the system to execute. In this situation, the following competency questions arise:

- If I am going a certain speed in specific terrain and I see an object of a particular type, what are the consequences of running into it or of avoiding it?
- How will a collision affect my payload or my ability to complete my mission?

4.4 Our initial focus

As shown above, an ontology component promises to be helpful in many aspects of the 4D/RCS architecture. We have decided to focus our initial efforts on the value judgement and behaviour generation components, specifically in the area of assisting the planner in deciding the most cost-effective plan. We will use an ontology component to help the planner determine the consequences of colliding with another object. The rationale for choosing this is as follows:

- Planning systems, in general, have a need to reason over information in the environment. Many planning systems are also at a level of maturity such that they are ready to incorporate ontologies into their framework.
- Recent advances in planning technologies have focused on building planning graphs based on attributes of objects in the environment (Balakirsky, 2003). Ontologies are well positioned to play a key role in allowing the planner to better understand the environment to be able to build more appropriate planning graphs.




Figure 2 Simple Driving Scenario

- Cost functions for collision consequences must include factors beyond those commonly used in kinematics or qualitative physics models, such as prohibitions on injuring humans and animals. Such factors are easy to accommodate in an ontology.
- Assigning costs to the consequences of collisions lends itself well to being tested in simulated environments. In this approach, we are not dependent on true sensor data for testing; we can use simulated objects to test the approach. Once the approach is well tested, we can then move the code to the actual vehicle.
- The approach allows us to easily tweak various parameters, such as the type of vehicle we are driving, the type of object in the environment, the speed we are travelling at etc., thus allowing us to show varied results for different situations.
- The classification hierarchy enables general rules for collision damage to be stated which cover a wide variety of types of vehicle and object. This is an advantage over having to input many specific rules.

The remainder of this paper will describe the details of our initial work on researching this approach.

5 Scenarios

In its full generality, the problem of automated vehicle navigation is extremely challenging. Our current focus is navigation on a roadway. We start with the simple scenario illustrated in Figure 2. Our vehicle (labeled OV) is in the left lane of a four-lane two-way undivided highway. An object is detected in our lane. The goal is to formulate an optimal route plan that takes into account the potential damage from a collision with the object. The main role of the ontology component is (initially) to provide assessments of collision damage. We will take into account not only damage to the vehicle, but also damage to the payload and to the object itself. This information is used to plan a route that either goes around the object or collides with it.

A number of parameters may be varied in this scenario. These include the type of vehicle we are controlling, the speed at which we are travelling, the payload we are carrying, and the type of object in our path that may be an obstacle. We will show how varying these parameters will change the plan that is ultimately executed. For example, if the object is a newspaper in the middle of the roadway, then

- the ontology component will conclude that no damage will occur,
- the value judgement component should decide that colliding with it has no real cost and
- the behaviour generation component will conclude that the best course of action is to maintain the current lane (because changing lanes always accumulates additional risk over maintaining your lane).

However, if the object were a large cinder block, significant damage would be likely and the final route should be quite different.

Deciding on a good route to follow is the job of the behaviour generation component. It generates plan segments whose costs are evaluated by the value judgement component. Each cost




Figure 3 Data Flow

is a function of a number of factors, one of which is the cost of running into objects on the roadway. The behaviour generation component identifies potential object-vehicle collisions in plan segments and passes these to the ontology component, which determines the likely damage to be incurred. The value judgement component must then find the best series of plan segments that achieve the system's overall goal.

The ontology component is equipped with knowledge about many kinds of vehicle and object, and the kind of damage that can arise from different collisions. This is used to determine the damage that would be caused by a collision. This information is returned to the value judgement component. In turn, the value judgement component uses this information to assign a cost to the plan segment. In the future, this operation may also be performed in the ontology component.

After establishing the basic concept with this simple scenario, we will progressively add more complexity and realism to this and other related scenarios. Note that the introduction of other traffic requires the autonomous vehicle to account for the location and speed of the other traffic when determining which path to take.

Figure 3 shows the data flow of information between behaviour generation, ontology and value judgement components. The vehicle type, object type, payload type, object position and traffic information are all knowledge that has been previously stored in the autonomous vehicle's world model, and that is fed to the behaviour generation component at planning time. When developing a plan, the behaviour generation component examines the current state of the system and simulates a set of discrete possible actions that the vehicle may take (accelerate, decelerate, change lanes etc.). As part of this simulation, the vehicle trajectory is examined in space-time to determine if a collision with an object takes place (Balakirsky, 2003). If a collision is expected to occur, the vehicle type, object type, payload type, vehicle closing speed and vehicle integrity are passed to the ontology for evaluation.

The ontology component takes all of the input and uses a reasoning engine to determine the expected damage due to a collision based upon its internal representation. The ontology component returns estimates of the integrity of the vehicle, the payload and the object collided with. These are passed to the value judgement component.

The value judgement component determines a single cost factor based on (1) the damage assessment from the ontology component, (2) previously defined mission objectives (e.g. get to a location as quickly as possible, get to a location as safely as possible etc.) and (3) simulation results from other knowledge layers. This cost factor is fed back to the behavior generation component where it is used in the construction of a planning graph. An A* search (Nilsson, 1998) is performed

on this planning graph which causes this cycle to repeat for all plan segments that need to be evaluated. When the graph search is complete, the path that results in the lowest cost is executed.

It is important to note in this example that there are many other factors that affect the cost of a path apart from collision with objects. Another category of such costs would be obeying the rules of the road. In the larger picture, all of these costs are evaluated to determine an overall cost for an action. It should be noted that the summation of these very different categories of cost into a single cost value is not a trivial matter. In the example described above, we are only looking at determining the consequences of colliding with other objects to show the value of the ontology. As we progress, we may expand the role of the ontology to include the overall value judgement evaluation as well.

6 Objects and obstacles

A major issue with route planning for autonomous vehicle navigation is to identify and determine the importance of potential obstacles in the vehicle's path. In the short term, our goal is to develop a knowledge base that can determine the extent to which a given object is an obstacle to a given vehicle in a given situation. The knowledge base will include an ontology of objects, vehicles and situations with associated inference rules. The rules determine the "degree of obstacleness", which is ultimately expressed in terms of a cost, and used in the value judgement component as described in the previous section. In the longer term, we aim for the ontology and inference rules to be the inspiration and foundation for a future general theory of objects and obstacles that can be reused in a wide variety of situations.

6.1 What is an obstacle?

What is an obstacle? The dictionary provides a good definition for our purposes: an obstacle is "something that impedes progress or achievement".² However, being an obstacle is not an inherent property of an object – a crate of oranges is a major obstacle to a moped, but may be no obstacle to a tank. Contrast this with the property of being a person. You cannot be a person one day and not a person the next day unless you cease to exist entirely – it is an inherent property of its instances.

An obstacle is a role that an object plays in a certain situation. A person may not be an obstacle to a certain vehicle if s/he is sitting in their living room, but if that same person walks into the middle of the road a short distance in front of that vehicle, they become an obstacle. A general theory of obstacles could be deployed to define a set of conditions that must occur for a given object to take on the role of an obstacle to a certain vehicle. The determination of the extent to which an object is an obstacle (or not) depends on the *relationship* between the object and another entity (for us, an autonomous vehicle). If the relationship includes impeding the progress of the vehicle, or impeding the vehicle's ability to carry out its goals, then the object is an obstacle.

In Guarino and Welty's (2000) taxonomy of property types, being an obstacle is a role. This means two things. First, being an obstacle is an optional property for all its instances. This is an alternate way to express the fact that being an obstacle is not an inherent property of an object, as noted above. Second, being a role means that the very existence of every instance of obstacle depends on the existence of some *other* instance (i.e. the thing being impeded). Something cannot be an obstacle unless there is something else that is being impeded in some way.

6.2 Ontology of objects

We are developing a formal representation of different types of objects that we expect to encounter in various environments, along with their pertinent characteristics and relationships with other objects. Initially we are focusing on on-road driving, so we are representing categories of object such as other vehicles, pedestrians, animals, debris, speed bumps etc. Each one of the objects that

² Merriam Webster online dictionary: <http://www.m-w.com/cgi-bin/dictionary>.

fall under these categories has a set of characteristics that describe them and help us to understand the damage that may be caused by colliding with them. For example, a certain type of debris may have a set of dimensions, a weight, a density, a velocity etc.

The ontology and its associated reasoning engine should be able to provide, as an output, a damage assessment in the event of a collision between our vehicle and a given object, based upon

- the type of autonomous vehicle,
- the type of object we are colliding with,
- the closing speed of our vehicle with the object and
- the integrity of our vehicle, i.e. what damage has already occurred to our vehicle, if any.

Based on this information, the ontology will provide a damage classification pertaining to

- the vehicle's integrity (initially only assigning damage to the bumper, wheels and overall vehicle, but will eventually include other components of the vehicle),
- the obstacle's integrity and
- the vehicle payload's integrity

During the initial stages of this work, we are assuming that the object in the environment is stationary. We are also assuming that all collisions are head-on collisions. This assumption eliminates the need to account for the angle of collision.

In order to provide the damage classifications, the expressiveness of the ontology must be such that it is able to represent concepts such as

- the type of vehicle that is being autonomously controlled and its pertinent characteristics,
- the objects that are being encountered in the environment and their pertinent characteristics,
- the payloads that the vehicle is carrying and their pertinent characteristics
- severity classifications of damage,
- damage types,
- terrain information (initially fixed as paved roads) and
- collisions (e.g. a certain type of vehicle with a certain type of object).

6.3 Damage assessment categories

As mentioned in Section 5, the main output of the ontology component will be the collision damage assessments of the vehicle, the object that it is colliding with and the payload. The damage assessment for each item will be expressed as one of five levels of severity: None, Minor, Moderate, Severe and Catastrophic. Table 1 defines the meaning for each level as it applies to the vehicle, the object and the payload. Note that the levels of severity of damage are inverse measures of the integrity of the given item. That is, a low amount of damage corresponds to a high integrity, and vice versa.

6.4 Estimating collision damage

So far we have introduced the general idea of an obstacle and specified some requirements for an ontology of objects and their characteristics in the roadway driving domain. Initially we have chosen to model the notion of an obstacle as costs expressed in terms of collision damage to the vehicle, the payload and the object collided with. In the future we will add other costs such as the extent to which the vehicle's goals are put at risk, and the priorities of those goals. These costs are used by the behaviour generation component in creating an optimal route plan. With all these elements in place, we will now consider how to estimate what the actual collision damage will be in a variety of situations.

There are many approaches that could be used to make this estimation. These include

- numerical simulation tools which model the physics of weight, materials, shapes, density, momentum etc. to compute impact damage,

Table 1 Meaning of levels of severity

	Vehicle	Object	Payload
None	No damage to vehicle	No damage to object	No damage to payload
Minor	Damage to vehicle will not affect vehicle performance	Damage to object will not affect object overall integrity	Damage to vehicle will not affect payload
Moderate	Moderate probability of vehicle damage, maintenance required	Damage to object will affect object integrity, but will not result in object destruction	Moderate probability of payload loss
Severe	Major loss of functionality/Integrity of vehicle likely	Major destruction of object	Major payload loss
Catastrophic	Vehicle loss	Object destruction	Payload loss

- probabilistic models,
- fuzzy logic and
- symbolic logic.

Our position is that no one of these techniques is likely to be adequate in all circumstances. In the longer term, we hope to explore all of these approaches, determining when each is most suitable, and to discover ways to use different approaches in concert to achieve superior results. Our current work focuses on the symbolic-logic approach. Our hypothesis is that even when logic-based inference is not sufficient, the core ontology of objects and characteristics will remain useful as a conceptualisation and vocabulary for expressing rules and procedures for estimating damage.

For our initial experiments, we have constructed a small ontology using OilEd (Bechhofer *et al.*, 2001). Using a description logic (Baader *et al.*, 2002) tool has two advantages for us. First, the classifier detects logical errors in the ontology, which greatly increases confidence that the ontology is correct. Second, it is very fast at doing inference. This is important because the planner needs to query the ontology component up to a few hundred times a second to get damage estimates for the many nodes being explored in the search space.

Classification reasoning works for our initial problem characterisation because we have identified categories of damage severity. We use it to answer the following three questions:

- In this situation, what is the likely damage that will occur to the vehicle?
- ... to the payload?
- ... to the object being collided with?

We defined a class called *Situation* which has various characteristics or attributes, each modelled by functional relations with *Situation* as the domain. The key characteristics of a *Situation* that will determine the damage classification are the vehicle, the payload and the object which the vehicle may collide with. These functional relations are called *hasVehicle*, *hasPayload* and *hasPotentialObstacle*, respectively. We also have attributes that are used to define the damage categories in Table 1. For example, the class *VehicleIntegrityMinor* is defined to be the class of all *Situations* such that the value of the functional relation *hasVehicleIntegrity* attribute is *Minor*.

We have a simple taxonomy of physical objects including various types of vehicle, and other objects such as bricks, newspapers etc. that may be in the vehicle's environment. These objects have characteristics such as weight, speed, density etc. that are important in determining the damage category. Initially we created some qualitative categories for measuring these characteristics, such as low, medium and high for weight or density.

Finally we created some axioms which specify how to classify a given situation in terms of the categories in Table 1. Here is a simple example:

A *Situation* such that

- the value of the *hasPotentialObstacle* relation is restricted to be of type *SmallDenseObject*

&

- the value of the *hasVehicle* relation is restricted to be of type *Car* is a subclass of *VehicleIntegrityModerate*.

We created some fictitious situations to test these axioms. For example, the situation whose *hasPotentialObstacle* relation is a brick and whose *hasVehicle* relation is a Toyota Corolla will be classified by this rule under *VehicleIntegrityModerate*. This is inferable because a brick is a *SmallDenseObject* (by virtue of its weight and size), and a Toyota Corolla is a subclass of *Car*.

From our initial tests, it is clear that there are limits to using a description logic reasoner for our task. For example, we cannot directly express rules that conclude that Integrity of the Vehicle and the Payload is unchanged, or that it should be incremented by one level in the damage severity scale. There are also limitations with respect to concrete domains that may be needed to reason with numbers. Similar problems were discovered in an attempt to use description logic classification to implement a semantic publish and subscribe system (Uschold *et al.*, 2003).

We are exploring alternatives for solving this problem. A variety of workarounds are possible. Also, it may be that description logic classification will only play a limited role in reasoning about obstacles. More general techniques may be required, such as production rules, probability or fuzzy logic.

This is a first pass, and as the project progresses the damage assessment categories in Table 1 will evolve into a complex multi-dimensional analytic device. There are several issues that are relevant: functionality and integrity of the entities are two key ones. Factored across these two issues is the likelihood of damage occurring. There may be many more categories and sub-categories that emerge from our analysis. Also, in addition to assigning a damage assessment to the overall vehicle, we will assign individual damage assessments to parts of the vehicle, such as the wheels and bumpers. We will then use these individual damage assessments to infer the damage assessment that should be assigned to the overall vehicle.

7 Future work and conclusion

The overall goal of this work is to apply ontologies to enhance the capabilities and performance of autonomous vehicles, particularly in the area of navigation planning. In order to do this, we are initially using an ontology to determine the damage resulting from collisions between autonomous vehicles and different types of object that could be encountered during on-road driving.

The scenarios presented in this paper will be implemented in a simulated environment and are scheduled to be completed by December 2003. In this evaluation system, we will include multiple types of vehicle, object and speed of collision. We hope to show how the plan that is generated, based on different combinations of input, changes as a function of the expected damage due to collision.

In the future, we plan to extend this work to find answers to questions such as:

- What is the nature of a “theory of obstacles”? How will it be integrated with the ontology of objects in the vehicle’s environment?
- What existing general theories and formal ontologies can be leveraged to create a theory of obstacles?
- To what extent can a general theory of obstacles be adapted to a wide variety of autonomous vehicle applications? Can we have a single ontology for multiple types of vehicle and context?

How much will they have to be tailored? This is analogous to the long-term question about Standard Upper Ontologies (SUOs), but within a limited domain. Can there be an SUO of obstacles?

- How can symbolic reasoning methods be used in conjunction with other approaches, such as probabilistic reasoning, fuzzy logic and numerical simulation?
- How can ontologies be linked to other types of representation, including sensor data, and other techniques for object identification (e.g. data and information fusion)?
- How can we leverage and/or complement a recent effort on applying ontologies for data fusion with the work described here on using ontologies for autonomous vehicle navigation? Attendees at a recent workshop on this topic provisionally agreed that: “good ontologies yield good fusion systems” (Llinas & Little, 2002). One obvious area of overlap is the object identification task in data fusion.
- Will the response times for ontology reasoning be fast enough to be useful in a real-time environment?
- What will be the best mechanisms for ontology sharing among different autonomous vehicles?
- Can using formal ontologies increase the possibility of having different autonomous vehicles be able to communicate among one another with reduced ambiguity? This would be particularly useful where multiple vehicles may be working towards a common goal.
- Can semantic integration techniques using ontologies be leveraged with multiple heterogeneous autonomous vehicles working together?
- How can one evaluate the performance of the ontology? Where does the ontology really add leverage compared to approaches not using ontologies? For example, does the ontology really help increase the ability to deal with dynamically changing environments? Which approaches are better in what circumstances?

The initial scenarios discussed in this paper are relatively simplistic, to allow us to focus on the technical issues surrounding implementing ontologies in autonomous vehicles. As the technical details are resolved, our scenarios will increase in complexity, and will include moving obstacles and will also account for angles of collision. We will also look at how to apply ontologies to other aspects of autonomous vehicle operations and will test the approaches developed in actual vehicles.

References

- Albus, J, 1991, “Outline for a theory of intelligence,” *IEEE Transactions on Systems Man and Cybernetics* **21** 473–509.
- Albus, J, 1995, “The NIST Real-time Control System (RCS): an application survey” *Proceedings of the AAAI Spring Symposium Series*.
- Albus, J, 2002, “4D/RCS a reference model architecture for intelligent unmanned ground vehicles” *Proceedings of SPIE Aerosense Conference* Vol. 4715.
- Albus, J *et al.*, 2002, “4D/RCS Version 2.0: A Reference Model Architecture for Unmanned Vehicle Systems,” NISTIR 6910, National Institute of Standards and Technology, Gaithersburg, MD.
- Albus, J and Meystel, A, 2001, *Engineering of Mind* John Wiley & Sons, Inc.
- Arkin, RC, 1989, “Towards the unification of navigational planning and reactive control”.
- Baader, F, McGuinness, D, Nardi, D and Patel-Schneider, F, 2002, *Description Logic Handbook: Theory, Implementation and Application* Cambridge University Press.
- Balakirsky, SB, 2003, *A Framework for Planning with Incrementally Created Graphs in Attributed Problem Spaces* Akademische Verlagsgesellschaft Aka GmbH.
- Bechhofer, S, Horrocks, I, Goble, C and Stevens, R, 2001, “OilEd: a reason-able ontology for the semantic web” *Proceedings of the Joint German Austrian Conference on AI* 396–408.
- Brooks, RA, 1985, “A robust layered control system for a mobile robot” MIT AI Lab, AI Memo 864.
- Donlon, JJ and Forbus, KD, 2003, “Using a geographic information system for qualitative spatial reasoning about trafficability” *Proceedings of the Qualitative Reasoning Workshop*.
- Forbus, K, Mahoney, JV and Dill, K, 2001, “How qualitative spatial reasoning can improve strategy games AIs: a preliminary report” *Proceedings of the 15th International Workshop on Qualitative Reasoning (QR01)*.

- Gazi, V, Moore, ML, Passino, KM, Shackleford, WP, Proctor, FM and Albus, JS, 2001, *The RCS Handbook* Wiley-Interscience.
- Gruninger, M and Fox, MS, 1994, "The role of competency questions in enterprise engineering" *IFIP WG 5.7 Workshop on Benchmarking Theory and Practice*.
- Guarino, N and Welty, C, 2000, "A formal ontology of properties" LADSEB/CNR Technical Report 01/2000.
- Llinas, J and Little, E, 2002, "An ontology action plan for the information fusion community: results of a DARPA/CMIF workshop".
- Malyyankar, RM, 1999, "Creating a navigation ontology" *Proceedings of the Workshop on Ontology Management, AAAI-99*.
- Meystel, A, Pouchard, L, Dawidowicz, E, Welsch, L, Schlenoff, C and Gruninger, M, 2002, "Issues raised by the panelists of the 'What is the role of ontologies in performance evaluation?' Panel" *Proceedings of the PerMIS '02 Conference*.
- Nilsson, NJ, 1998, *Artificial Intelligence: A New Synthesis* Morgan Kaufmann Publishers, Inc.
- Schlenoff, C, "Linking sensed images to an ontology of obstacles to aid in autonomous driving" *Proceedings of the 18th National Conference on Artificial Intelligence: Workshop on Ontologies for the Semantic Web*.
- Uschold, M, Clark, P, Dickey, F, Fung, C, Smith, S, Uczekaj, S, Wilke, M, Bechhofer, S and Horrocks, I, 2003, "A semantic infosphere" *Proceedings of the 2003 International Semantic Web Conference (ISWC 2003)*.