Ontologies for World Modeling in Autonomous Vehicles

Mike Uschold, Ron Provine, Scott Smith
The Boeing Company
P.O. Box 3707,m/s 7L-40
Seattle, WA USA
98124-2207
michael.f.uschold@boeing.com

Craig Schlenoff, Stephen Balikirsky
NIST
Maryland
schlenof@cme.nist.gov

Abstract

We are in the initial stages of a collaboration between Boeing and NIST. We are exploring the hypothesis that it is beneficial to use ontologies to augment traditional world modeling technologies for autonomous vehicles. Our approach is to develop a theory of obstacles represented as an ontology. It will provide the basis for identifying and reasoning about potential obstacles in the vehicle environment in order to support navigation. We will develop a prototype implementation that incorporates the obstacle ontology and an associated reasoner into an existing autonomous system infrastructure. This infrastructure is based on the 4D/RCS architecture developed at NIST.

1 Introduction

We are in the initial stages of a project that is testing the hypothesis that ontologies can provide benefits in the context of autonomous vehicle navigation. This is a collaboration that leverages and applies ontology expertise at the Boeing Company to an existing autonomous vehicle effort at the National Institute of Standards (NIST). In order to get early feedback from the community, we describe our plans for this collaboration and our progress to date.

A major challenge in autonomous vehicle navigation is the ability to maintain an accurate representation of pertinent information about the environment in which it operates. The inability to do this well hinders effective task planning and execution, especially navigation. Efforts on-going at NIST are applying the 4D/RCS reference model architecture [Albus, J. et.al. 2002] to control an autonomous High Mobility Multipurpose Wheeled Vehicle (HMMWV). An explicit component of the 4D/RCS architecture is a world model which represents the vehicle environment (see figure 1). While the need for ontologies for world modeling is acknowledged, it has not been addressed.

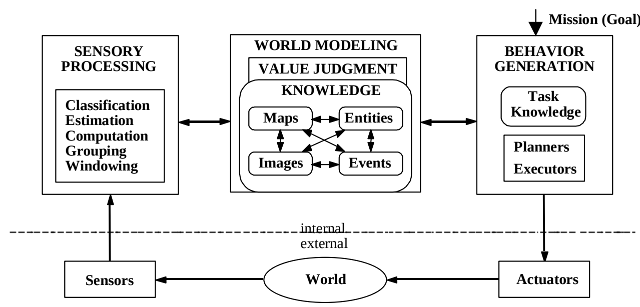

Figure 1. *The basic internal structure of a 4D/RCS control loop. Sensory processing performs the functions of windowing, grouping, computation, estimation, and classification on input from sensors. World modeling maintains knowledge in the form of images, maps, entities, and events with states, attributes, and values. Relationships between images, maps, entities, and events are defined by pointers. These relationships include class membership, ontologies, situations, and inheritance. Value judgment provides criteria for decision making. Behavior generation is responsible for planning and execution of behaviors.*

The overall goal of this work is to explore the hypothesis that ontologies can play a significant role in enhancing the capabilities and performance of autonomous vehicles, particularly in the area of navigation planning. To support navigation, an ontology needs to include:

- various objects that an autonomous vehicle is expected to encounter in its environment, and their important characteristics;
- factors that affect the motion of objects, for example: *obstacles*, road networks, rules of the road;
- actions that an autonomous vehicle is able to perform.

By introducing an ontology (or set of ontologies) into an autonomous vehicle's knowledge base, we can achieve many potential benefits. One is the potential for reuse and modularity. A general theory of obstacles, for example could apply in a broad range of autonomous vehicles, adapted to the special circumstances of each. Also, for a given vehicle system, an ontology provides the opportunity for a more centralized approach for representing and reasoning with information about the environment. Different modules could query the ontology, rather than having different pieces of the problem scattered across different modules. This has a corresponding benefit in cheaper and more reliable maintenance. An ontology can also extend the range of important questions that can be answered to support navigation planning. For example:

- Based upon sensor data, what are the objects we perceive in the environment at a given time?
- To what extent is a particular object a potential obstacle?
- What is our risk of colliding with the object assuming the motion patterns do not change?
- What are the appropriate actions in a given situation?

Finally, there is potential for increased flexiblity of response for the autonomous vehicle. Methods that rely on pre-classification of certain kinds of terrain in terms of their traversability [Donlon & Forbus 1999; Malyankar 1999] are important, but do not support reasoning with obstacles in a more dynamic context.

This effort will focus on assisting in vehicle navigation. For successful navigation planning, an autonomous vehicle is required to know the extent to which a given object may impede its progress. We will develop a *theory of obstacles* represented as an ontology to determine this for a variety of objects, vehicles and situations. Necessarily, it will be tightly integrated with the ontology of objects in the environment. This will complement other work at NIST that is addressing the representation of rules of the road and road networks.

2 Approach

Our primary focus is the role of obstacles in navigation planning. A long term aim is to develop a comprehensive and reusable ontology of obstacles that can be used in a wide variety of contexts, different vehicle types and environments. We will implement a proof of concept demonstrator which incorporates an obstacle ontology of modest scope with an associated inference engine into an existing autonomous vehicle infrastructure currently being developed at NIST. We will use a simulation tool to generate object data that would otherwise be obtained from processing sensor data from an actual vehicle.

The work will proceed in three phases.

1. Identify the requirements;
2. Create an obstacle ontology related to the objects in the vehicle environment;
3. Implement the proof of concept demonstrator.

We will use a scenario-driven spiral development method starting with a small set of initial requirements, and then repeating the steps adding new requirements and functionality. We now elaborate on these steps, summarizing our progress to date.

3. Identify Requirements

The requirements phase involves the following:

1. identify a scenario;
2. identify competency questions;
3. scope the ontology;
4. identify the representation and inference requirements.

3.1 Scenario

First we identify a scenario which demonstrates the utility of reasoning with an ontology of obstacles. For example, a simple initial scenario could involve a single vehicle driving down a road. We will consider different objects that may be on the road, e.g. small cardboard box or a crate of oranges. We will also consider different traffic conditions. The appropriate action with a small cardboard box in the vehicle’s lane is to drive around it. The same situation in heavy traffic, might require going over the box which will is unlikely to cause damage to the vehicle. The same scenario with crate of oranges is more complex. The risk of damaging the vehicle by running into the object must be balanced by the risk of an accident causing damage to one or more vehicles on the road. If the other vehicles are driving in a predicable manner, it may be safe to swerve to avoid the object, if there is a nearby vehicle that is driving erratically, it will be less safe. The navigation planner will view this as a different costs presented by the existence of the obstacle.

In later cycles of the spiral, we will elaborate on the initial scenario, and/or identify a set of related scenarios that could affect the context-specific characteristics of an obstacle (e.g., the speed of a vehicle could affect the damage that could be done by colliding with the obstacle).

3.2 Competency Questions

Next we identify specific competency questions [Gruninger and Fox 1994] that the reasoner must answer using the ontology to support the scenario. In the initial scenario, there would be a small number of simple questions. Here we include a broader range of questions that might come up in more complex scenarios.

1. 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 sensor input and other reasoning?
   e. how do I determine that confidence?
2. What other information do I know about the object once I identify it? Does it have ammunition and am I in it's range, is it friend or foe, etc.
3. If I seen a object with certain properties and I'm not sure what it is, what additional information should I gather so that I will be better able to identify the object? This information could be used to task sensors for gathering further information.
4. If I am going a certain speed in specific terrain and I see an obstacle of a particular type, what is the cost of running into it, or of avoiding it?
5. 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).
6. 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?

   This could depend on context. If a man is standing holding a lollipop sign with slow/stop on either side, then he is unlikely to move into traffic.

3.3 Scope the Ontology

Next we identify the scope of the ontology that is required to answer the competency questions to support the identified scenario. The questions listed above are broader in scope than would be required to support the small initial scenario.

3.4 Representation and Inference

Finally, we will identify the representational and inference requirements needed to answer the identified set of competency questions. This will be the basis for selecting ontology development and inference tools for subsequent phases.

4 Create Ontology

The ontology creation phase involves the following:

1. Literature search on obstacle ontologies;
2. Select appropriate ontology representation, inference and development tools;
3. Create formal representation of obstacles and objects to meet requirements from phase 1;

4.1 Literature Search

We have begun to perform a literature survey to determine relevant work that can be leveraged in the development of a general theory of obstacles. Google returns no hits on obvious search patterns such as: "ontology of obstacles" or "obstacle ontology" (except for the authors' prior work). The pattern, "theory of obstacles" returns many and only false hits. When this is conjoined with "navigation" or "robot" or "autonomous" there are no hits at all. This is an indicator that the idea of having an explicit theory or ontology of obstacles for autonomous system navigation purposes may be relatively new.

The most closely related work we found is in the area of determining 'trafficability'. This is defined to be: "a measure of the capability for vehicular movement through some region" i.e. specific kinds of terrain [Donlon & Forbus 1999]. This work is being done in the context of traditional GIS algorithms that may be used for route planning. They are being augmented with qualitative reasoning techniques. 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 its traversability. Slope, hydrography, vegetation and other things are taken into account. For example, if the slope angle is greater than 45 degrees, this would be severely restricted for most 4-wheel vehicles.

Similar work is reported in [Malyankar 1999]. The creation of "navigation ontology" in a marine environment is discussed. It is also set in the context of GIS. These and other sources will be studied and mined for ideas that we hope to generalize and apply to create an ontology of obstacles. Neither work addresses the issue of reasoning about obstacles in real-time from sensor data, it is all based on pre-classifying known terain. These approaches therefore would not be able to handle our crate of oranges example. Such dynamic capability will be a focus of our research.

4.2 Select Ontology Tools

We will then select an appropriate ontology language, inference engine, and development tools. There is a wide variety of tools to select from. This will be performed by 1) analyzing and determining an appropriate formalism (or set of formalisms) in which to represent the ontology of obstacles, 2) analyzing and determining an appropriate formalism (or set of formalisms) for inference engines, 3) identifying suitable formalism/inference engine combinations, 4) selecting the best combination, and 5) selecting a development tool (e.g. OilEd, Protégé). This decision will also be affected by system requirements arising from the NIST software infrastructure.

4.3 Create Formal Representation

We consider two aspects of creating a formal representation of the ontology:

1. conceptual analysis
2. formalization

The first entails identifying the important objects and relationships and finding a way to think about obstacles and their relationship to objects. The second is to represent the results of this analysis and design in a formal language. We report here on some early analysis. We have not begun the formalization stage.

A theory of obstacles is different from an ontology of objects per se. An object may or may not be an obstacle, and this can change over time. One of the interesting questions of our project is: what is the relationship between a theory of obstacles and ontology of objects. Some of the factors that determine whether something is an obstacle are: the vehicle, the context, and to some extent the purpose or goals of the vehicle. The same object, say a small bush will be an obstacle for a small car, but not for an army tank. For a given vehicle, say a car, the same object may be an obstacle at high speed, but not at low speed. An object's location also determines the extent to which it will be an obstacle.

We distinguish two types of characteristics about objects:

- *static characteristics* - characteristics about object that are not a function of the context in which is it is viewed (dimensions, location, velocity, armed/not armed, color, etc.)
- *inferred characteristics* - characteristics that need to be determined through reasoning (is the object of importance?, is the object a threat?, etc.). This would be a function of context, intention, environment, etc.

We will also have to determine which characteristics will be represented in the ontology, and which will be represented outside of the ontology (e.g., cost of running into obstacles?).

5 Implement Prototype

We will implement a proof-of-concept scenario in which the planner develops a plan around/through obstacles based on the retrieved characteristics of the obstacles from the ontology. This entails integrating the ontology of obstacles with NIST's planner and simulation package. Initially, the simulation package will send the exact obstacle (object classification optional) that is being encountered to the ontology and the ontology is sending back the important characteristics of that object.

A cost model will be developed that represents how to respond to different obstacle characteristics. Also, we will ensure that all information provided by the ontology, and the associated inferences, are viewable by the user to allow for 'white box' planning and development.

6 Issues and Challenges

There are many open questions and technical challenges posed by this work. Some of these are listed below:

- 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?
- How can symbolic reasoning methods be used in conjunction with probabalistic reasoning for use in autonomous vehicle navigation?
- How can ontologies be linked to other types of representations, 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 and 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?
- 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 vehicles and contexts? How much will they have to be tailored? This is analogous to the long-time question about standard upper ontologies (SUO), but within a limited domain. Can their be a SUO of obstacles?
- What will be the best mechanisms for ontology sharing among different autonomous vehicles?
- Using formal ontologies increases 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 toward a common goal.
- Can semantic integration techniques using ontologies be leveraged with multiple heterogeneous autonomous vehicles working together?
- What other aspects of autonomous systems may ontologies add value besides navigation planning?
- How can one evaluate the performance of the ontology? Where does the ontology really add leverage compared to approaches not using ontologies? For ex-

ample, does the ontology really help increase the ability to deal with dynamically changing environments? When would these other approaches be preferred, and when would ontology-bases approaches be preferred?

References

[Albus, et.al. 2002], Albus, James S et al "4D/RCS Version 2.0: A Reference Model Architecture for Unmanned Vehicle Systems," NISTIR 6910, National Institute of Standards and Technology, Gaithersburg, MD, 2002.

[Albus & Meystel 2001] Albus, James S., Meystel, Alexander M. eds. 2001. Engineering of Mind: An Introduction to the Science of Intelligent Systems New York: John Wiley and Sons.

[Donlon & Forbus 1999] Donlon, J.J. and Forbus K.D. Using a Geographic Information System for Qualitative Spatial Reasoning about Trafficability. *Proceedings of the Qualitative Reasoning Workshop*. Loch Awe, Scotland. 1999. http://www.qrg.northwestern.edu/papers/files/Donlon_Forbus_QR99_Distribution.pdf

[Gruninger and Fox 1994] Gruninger, M., Fox, M.S., The Role of Competency Questions in Enterprise Engineering, IFIP WG 5.7 Workshop on Benchmarking Theory and Practice, Trondheim, Norway, 1994.

[Llinas and Little 2002] James Llinas and Eric Little: An Ontology Action Plan for the Information Fusion Community: Results of a DARPA/CMIF Workshop, November 2002. An http://www.infofusion.buffalo.edu/conferences_and_workshops/ontology_and_viz_ws/ws_products/ontology_action_plan/Ontology%20Action%20Plan.ppt

[Malyankar 1999] R. M. Malyankar: Creating a Navigation Ontology, Workshop on Ontology Management, AAAI-99, Orlando, FL, 1999. In Tech. Rep. WS-99-13, AAAI, Menlo Park, CA. http://www.eas.asu.edu/~gcss/papers/wk99.pdf

[Schlenoff, 2002a] Position statement at panel discussion on the "Role of Ontologies in Intelligent Systems" at Performance Metrics for Intelligent Systems (PerMIS) 2002 held at NIST in Washington, DC

[Schlenoff 2002b] Schlenoff, Craig “Linking Sensed Images to an Ontology of Obstacles to Aid in Autonomous Driving,” Papers of the AAAI Workshop Technical report WS-02-11, Proceedings of the AAI-02 Conference, Edmonton, Alberta, Canada, July 28-August 1 2002.