

Grounding Robot Sensory and Symbolic Information Using the Semantic Web

Christopher Stanton and Mary-Anne Williams

Innovation and Technology Research Laboratory, Faculty of Information Technology
University of Technology, Sydney, Australia
{cstanton, Mary-Anne}@it.uts.edu.au

Abstract. Robots interacting with other agents in dynamic environments require robust knowledge management capabilities if they are to communicate, learn and exhibit intelligent behaviour. Symbol grounding involves creating, and maintaining, the linkages between internal symbols used for decision making with the real world phenomena to which those symbols refer. We implement grounding using ontologies designed for the Semantic Web. We use SONY AIBO robots and the robot soccer domain to illustrate our approach. Ontologies can provide an important bridge between the perceptual level and the symbolic level and in so doing they can be used to ground sensory information. A major advantage of using ontologies to ground sensory and symbolic information is that they enhance interoperability, knowledge sharing, knowledge reuse and communication between agents. Once objects are grounded in ontologies, Semantic Web technologies can be used to access, build, derive, and manage robot knowledge.

1 Introduction

Robots interacting with other agents in complex dynamic physical environments require sophisticated and robust concept and knowledge management capabilities if they are to solve problems, communicate, learn and exhibit intelligent behaviour [4]. The Semantic Web is touted as the next stage in the World Wide Web's evolution, and involves the use of ontologies - collections of information that formally define the relations among terms. The structured nature of ontologies allows agents, such as autonomous robots, to make sense of the knowledge available on this global network.

For autonomous agents, ontologies can provide an important bridge between the perceptual level and the symbolic level, and in so doing they can be used to ground sensory information. Ontologies can be used to help link sensory data with symbolic representations by defining the associated sensory phenomena of real-world objects. Symbolic information can then be derived for the purpose of planning or communication. A major advantage of using ontologies to ground sensory and symbolic information is that they can be used to enhance interoperability, knowledge sharing, knowledge reuse and communication between agents. Once objects are grounded in ontologies, Semantic Web technologies can be used to access and manage robot knowledge.

Knowledge management requires the formation and evolution of concepts and categories. Categorization forms the foundation of intelligent behaviour. Robots, for example, need to categorize both physical and abstract entities of interest in their environment to achieve their goals. Ontologies are a practical tool for developing a conceptual network that robots can use to classify and identify objects with a high degree of exception tolerance. The ability to classify new information is essential for reasoning, problem solving, communication and learning [4].

OWL is a semantic markup language for Web resources that can be used to define ontologies (<http://www.w3.org/>). Using OWL, robots can build rich and grounded world models from a wide variety of internal and external knowledge (re)sources, such as sensors, ontologies, databases, knowledge bases, the Semantic Web, web services, and other agents.

In this paper we describe a Robot Soccer Ontology (RSO) in OWL. This ontology enhances the knowledge management capability of a robot and allows it to integrate information from multiple sensory and symbolic sources as well as understand its environment to the extent of communicating with humans and other agents. We use AIBO robots that play robot soccer to illustrate our framework and approach. Our area of concern is modeling how robots utilize concepts and knowledge as a result of their interaction with heterogeneous information landscapes. One of our motivations for this work is to build a robot soccer multiagent system in which each robot *knows* it is playing soccer and understands all the important elements of the game such as it is a ballgame with a set of rules that govern how it is played.

2 Our Application – Robot Soccer

The AIBO ERS-210 is an entertainment robot produced by SONY. The AIBO has a MIPS CPU with a clock cycle of 192 MHz and a 32 megabyte main memory. It is equipped with an array of touch sensors, on positions such as the head, back, chin and paws. The AIBO's vision system is driven by a single CMOS camera, while there is also an infrared range finder. All four legs, as well as the head, have three degrees of freedom allowing for a large range of motion. Sony's freely available OPEN-R software development kit provides a set of application programming interface library files and related tools that allow both querying and control of the AIBO's range of hardware through C++ code. The ERS-210 has wireless LAN capabilities that allow inspection of the robot's internal state and access to the vast resources of the Internet.

The SONY Four-Legged RoboCup Competition is held annually. Each robot soccer team is comprised of four AIBOs; of which one is designated as the goalkeeper. Most of the robot functions and soccer behaviours are driven by visual sensory information. The soccer field and other objects of importance in the game are uniquely colour-coded. The field is *green* in colour and defined by six bi-coloured beacons; one in each corner, and two on either side of the half-way line. One goal is *blue* and the other is *yellow*. The ball is *orange*. The beacons, goals and ball are all of known size. The beacons are used by the robots to determine their location on the field. Since

AIBOs do not possess stereo vision they cannot make independent and reliable judgments about distances to objects using vision alone, so algorithms based on vision data determine distance using known dimensions of objects. The vision processing system of the robot converts raw camera images in YUV format into conceptual information about physical objects. Pixels in the initial YUV image are classified using a set of colour concepts which are of value to the robots' task of playing soccer e.g. recognising the ball, the beacons, and the goals. Classified pixels are accumulated into larger "blobs" of colour. The blobs are then attributed to specific objects, e.g. beacons.

3 Grounding Sensory and Symbolic Information

The symbol grounding problem refers to the task of linking the internal concepts used by an agent in its world model for reasoning and decision making with the real world phenomena to which they refer [7]. An agent using ungrounded information could be considered delusional because it can perceive things that do not exist while failing to perceive things that do exist. A special case of the symbol grounding problem is anchoring [3]. Anchoring involves maintaining the link between an agent's internal representation, and the physical objects represented by these symbols. For example, in the RoboCup domain, grounding (or anchoring) could involve linking an orange blob in a robot's camera image with a soccer ball.

Simple systems that are purely reactive, such as an elevator, do not require a symbolic model of their environment, and thus avoid the grounding/anchoring problem entirely. However, for more complex systems a symbolic system can allow an agent to construct an internal model of the world, which importantly allows the agent to reason, plan and predict future states - activities integral to intelligent behaviour. Any agent functioning in a physical environment that has a symbolic reasoning component must have a solution for grounding or anchoring symbols. However, such solutions tend to be restricted to the particular domain for which the agent operates. There is a need for more general, less domain-specific solutions.

Harnard [7] argues that symbolic representations must be directly grounded bottom-up using iconic and categorical representations, where both iconic and categorical representations are non-symbolic. Iconic representations can be thought of as copies of sensory projections that preserve relevant features of the projection, where as categorical representations are reduced to the features of the projection that can determine category membership. By building categories that filter features of iconic representations, within category similarity differences compress, while between-category similarity distances expand, allowing for reliable classification of category membership.

The use of categories restricts the symbol grounding problem to a reduced set of elementary symbols, where each of these elementary symbols has a corresponding real world grounded object or event [1]. High-order symbols can be built through combinations of elementary symbols, making it possible to learn new concepts from pure linguistic definition. For example, the concept of zebra could be defined as

“horse + stripes”, with the terms “horse” and “stripes” grounded in the real world. By grounding bottom-up high order symbolic representations can be developed from elementary symbols.

Ontologies can provide an important bridge between the low-level, sub-symbolic Perception Layer, and high-level symbolic representations in the Deliberative Layer, by defining the properties of physical objects in terms of sensory data and events. In providing libraries of structured knowledge universally accessible to machines (including autonomous robots), we can move towards a generalised solution to the grounding problem. For instance, a robot confronted with a banana for the first time could use the semantic web to discover that this yellow, curved shaped object is a non-toxic fruit, and can not move on its own,. As such, ontologies represent a general domain independent mechanism for grounding sensory and symbolic information in autonomous robot systems.

4 Concepts and Knowledge Management on the Semantic Web

Categorization involves partitioning objects into cognitively useful groups, with these groups referred to as categories or concepts. Resultant concepts can be used to build knowledge; concepts can become predicates which are used to describe the world and its behaviour in a knowledge base, e.g. Situation Calculus [8].

Concepts help robots reduce the complexity of the information they need to acquire and manage. Furthermore, concepts give rise to broad powers of explanation. For example, without concepts robots would not be capable of representing visual information beyond the pixel level, and as a result they would not develop a world model that could support even simple forms of object recognition and reasoning. The ability to form and manipulate concepts explicitly enhances robots capacity for problem solving, communication, collaboration etc, as they forage around information rich heterogeneous environments.

We expect robots to respond appropriately to information acquired through their sensory systems. The ability to categorize new sensory information and to anchor it to physical objects in the real world allows a robot to behave sensibly in previously unencountered situations. Ontologies provide a powerful and practical tool for developing feature based categorizations of physical and abstract concepts.

According to Guarino [6] an ontology is a philosophical discipline, a branch of philosophy that deals with the nature and the organisation of reality, and for Aristotle it is the science of being. In the enterprise of building artificial agents, an ontology is an explicit specification of a conceptualization which involves terms and relations in a domain [5], or a shared understanding of some domain [9].

Ontologies are collections of information that formally define the relations among terms. Communication relies on the ability to share meaning – this can be achieved via ontologies. For example, two robot soccer knowledge bases may use different identifiers for what is the same concept e.g. one might use *beacon* and whilst the other uses *marker*. An agent that wants to compare or combine information across these databases must know that the two terms mean the same thing. Ideally, an agent

should have the capability to discover common/shared meanings for whatever database it encounters.

The most typical kind of Ontology for the web has a taxonomy and a set of inference rules. The taxonomy defines the class of objects and relations among them (see Appendix A for our Robot Soccer object model). Ontologies can be used for multiple purposes, from improving the accuracy of web searches to more advanced applications, such as Robot Soccer which can use ontologies to relate the information from a robot's sensors or knowledge base to the associated knowledge structures and inference rules. A markup language like OWL helps to develop programs that can tackle complicated problems whose answers require the fusion of information from multiple sources.

We describe our Robot Soccer Ontology in Appendix A. It is an explicit and formal description of the concepts (and categories of objects) in the domain of robot soccer and is composed of an object model and axioms. The concepts of interest are classes, where the properties of each concept describe important features and attributes of the concept, and restrictions on properties constrain roles of the properties. The ontology, a set of individual instances of classes, together with a set of axioms, constitutes our Robot Soccer Knowledge Base.

Classes describe concepts in the domain, and are the base objects of an ontology. In OWL each class is represented by a URI, and can be defined as a sub-class of an existing class. A description of the class can be added, as can restrictions to that class. These restrictions are based on properties. For example, the super-class **Thing** represents all objects of interest in Robot Soccer for the purpose of our application. Specific objects are instances of **Thing**, e.g., a goalie would be an instance of **Goal Keeper**, a subclass of **Player** which, in turn, is a subclass of **Thing**, while a cardinality restriction states that there is only one **Goal Keeper** per **Team**.

Axioms are used to state facts and rules about classes and their relationships. Axioms are provided in addition to class definitions which are essentially restricted forms of subsumption/equivalence axioms. OWL allows class definitions to assert the disjointness or equivalence of classes. An example of an axiom is that for all teams, there exists only one defensive player who may enter the penalty area, and that player is the designated goal keeper. Other ontologies can be incorporated into OWL using namespaces which are URIs that have been imported from other ontologies.

5 The Semantic Web and Web Services for Robots

The aim of the Semantic Web is to bring meaningful content to the World Wide Web, creating an environment where agents interact and accomplish sophisticated tasks. The Semantic Web is designed to offer access to ontologies, knowledge resources, and knowledge management tools and services. It allows people and agents, like robots, to embed meaning into web pages to enhance “understanding” of the content using the Resource Description Framework (RDF).

RDF represents data as triples; subject, predicate and its value. RDF Schema extend this binary propositional language's expressivity to that of a semantic net. RDF

Schema allow a wide variety of knowledge resources to be described. OWL is an example of a restricted semantic net that can be used to build complex ontologies. An OWL ontology is essentially a web page containing: (i) an optional owl:Ontology instance, (ii) a set of classes, (iii) a set of properties, and (iv) a set of restrictions relating the classes and properties. OWL extends RDF(S) by: (a) supporting XML Schema Datatypes rather than just string literals, (b) local restrictions, (c) enumerations, (d) class expressions, and (e) ontology and instance mapping.

Web services are self-contained, modular applications that have an explicit description of a service offered. They can be published, accessed and used on the web. Web services are currently based on a small set of XML based standards for message transfer (i.e. SOAP), Web Service description (WSDL), and Web Service discovery (e.g. UDDI framework). Web Services not in registries can be published using Web Service Inspection Language which assists search engines to find them. Logic for the Semantic Web can be added to web pages using a combination of OWL, Resource Description Framework (RDF), RDFS (RDF Schema) and XML.

We have developed a Robot Soccer Ontology which describes the domain of Robot Soccer for the RoboCup SONY 4-Legged League. The purpose of the RSO is to assist robots to play soccer by providing a mechanism for grounding their sensory and symbolic information to physical and abstract objects and concepts. In particular, the RSO is used to allow the robots to answer the following questions:

- Am I seeing a beacon? Which beacon am I seeing?
- Am I seeing a robot? What team is the robot I am seeing on?
- Am I seeing the goal? Which goal am I seeing?
- What is the score? Is my team winning?
- Which goal should I aim for? Where is the goal?
- Where are my team-mates?
- Which kick should I use?

The RSO can be implemented in OWL and placed on the internet for the robots to access. Once assessable it can be used for the robots to acquire and share information. Some examples of "Things" in the RSO represented in OWL are given below:

Robot Class Example.

```
<owl:Class rdf:ID="Robot">
  <rdfs:label> Robot</rdfs:label>
  <rdfs:comment>
    4-Legged League robots are on the red or blue team.
  </rdfs:comment>
  <rdfs:subClassOf rdf:resource="#ConcreteObject">
    <rdfs:subClassOf>
      <owl:Restriction owl:cardinality="1">
        <owl:onProperty rdf:resource="#hasTeam">
          </owl:Restriction>
        </rdfs:subClassOf>
      </owl:Restriction>
    </rdfs:subClassOf>
  </owl:Class>
```

Example Instance of Robot Class.

```
<Robot>
  <team> RedTeam </team>
  <position>Attacker</position>
</Robot>
```

6 Summary

Grounding involves the creation and maintenance of linkages between objects in a robots internal world model and objects in the environment. In this paper we described methods that can be used to implement the grounding of sensory and symbolic information embedded in a robots world model using ontologies designed for the Semantic Web. We use SONY AIBO robots and the robot soccer domain to illustrate our approach.

We developed a Robot Soccer Ontology and Knowledge Base that can be used to identify and reason about real-world objects and events. For example, the RSO defines the individual colours of the colour-coded RoboCup world in terms of sensory data, such as YUV pixel values.

A major advantage of using ontologies to ground sensory and symbolic information is that they enhance interoperability, knowledge sharing, knowledge reuse and communication between agents/robots. Once objects are grounded in ontologies, Semantic Web technologies can be used to access, build, infer from, and manage robot knowledge.

References

1. Cagenolsi, A., Greco, A., and Harnard, S., 2002, "Symbol grounding and the symbolic theft hypothesis", in A. Cangelosi and D. Parisi, (eds) *Simulating the Evolution of Language*, London: Springer.
2. Chalup, S., Creek, N., Freeston, L., Lovell, N., Marshall, J., Middleton, R., Murch, C., Quinlan, M., Shanks, G., Stanton, C., and Williams, M-A., "When NUbots Attack! The 2002 NUbots Team Report", Newcastle Robotics Laboratory Report, <http://robots.newcastle.edu.au/NUbotFinalReport.pdf>
3. Coradeschi, S., and Saffioti, A., 2000, "Anchoring symbols to sensor data: preliminary report", in *Proc. of the 17th National Conference on Artificial Intelligence (AAAI-2001)*, pp.129-135.
4. Gärdenfors, P. and Williams, M-A, *Building Rich and Grounded Robot World Models from Sensors and Knowledge Resources: A Conceptual Spaces Approach*, in the *Proceedings of the International Symposium on Autonomous Mini-robots for Research and Edutainment*, 123 – 133, 2003.
5. Gruber, T. R., 1993, "A translation approach to portable ontology specifications", *Knowledge Acquisition*, 5(2).
6. Guarino, N. Formal ontology, conceptual analysis and knowledge representation. *International Journal of Human-Computer Studies*, 43(5–6), 1995, pp. 625–640.
7. Harnard, S., 1990, "The symbol grounding problem", *Physica D*, Vol.42, pp.335-346.
8. McCarthy, J. and Hayes, P. J., 1969, "Some Philosophical Problems from the Standpoint of Artificial Intelligence", in D. Michie (ed), *Machine Intelligence 4*, New York: American Elsevier.
9. Uschold, M., and Grüninger, M., 1996. "Ontologies: Principles, methods and applications", *Knowledge Engineering Review*, 11(2).

Appendix A: Robot Soccer Ontology

