

Ontology Development Principles

Computers do not have the equivalent of a human cognitive system and therefore store data simply as the numbers and words that are entered into the computer. For a computer to interpret data it requires an information structure that provides at least some level of *context*. This can be accomplished utilizing an *ontology* of objects with characteristics and a rich set of relationships to create a virtual version of real world situations and provide the *context* within which *agent* logic can automatically operate.

This paper discusses the development of ontologies that serve to provide context for agents to interpret and reason about data changes in decision-support software tools, services and systems. The following brief explanation of key terms and concepts referred to in this paper is provided as an introduction for clarification purposes.

Ontology: The term ontology is loosely used to describe an information structure, rich in relationships that provides a virtual representation of some real world environment (e.g., the context of a problem situation such as the management of a transport corridor, the loading of a cargo ship, the coordination of a military theater, the design of a building, and so on). The elements of an ontology include objects and their characteristics, different kinds of relationships among objects, and the concept of inheritance. Ontologies are also commonly referred to as *object models*. However, strictly speaking the term ontology has a much broader definition. It actually refers to the entire knowledge in a particular field. In this sense an ontology would include both an object model and the software agents that are capable of reasoning about information within the *context* provided by the object model (since the agents utilize business rules that constitute some of the knowledge within a particular domain).

Information and context: Information refers to the combination of data with relationships to provide adequate *context* for the interpretation of the data. The richer the relationships the greater the *context* (i.e., meaning conveyed by the combination of data with relationships), and the more opportunity for automatic reasoning by software agents.

Information-centric: Software that incorporates an internal information model, such as an ontology, is often referred to as *information-centric* software. The information model is a virtual representation of the real world domain under consideration and is designed to provide adequate *context* for software agents (typically rule-based) to reason about the current state of the virtual environment. Since information-centric software has some understanding of what it is processing it normally contains tools rather than predefined solutions to predetermined problems. These tools are commonly software agents that collaborate with each other and the human user(s) to develop solutions to problems in near real-time, as they occur. Communication between information-centric applications is greatly facilitated since only the changes in information need to be transmitted. This is made possible by the fact that the object, its characteristics and its relationships are already known by the receiving application.

Agents: This term has been applied very loosely in recent years. There are several different kinds of agents. Symbolic reasoning agents are most commonly associated with

knowledge management systems. These agents may be described as software modules that are capable of reasoning about events (i.e., changes in data received from external sources or as the result of internal activities) within the context of the information contained in an internal information model (i.e., ontology). The agents collaborate with each other and the human users as they monitor, interpret, analyze, evaluate, and plan alternative courses of action.

Methodology for Developing an Ontology

The process of building an ontology consists of a number of steps to ensure the validity and completeness of the final product. The basic steps include:

Step (1): Identifying the purpose: The intended use of the ontology must be defined in some formal way to guide the development process. Without a well-defined purpose of the ontology, development can continue with no apparent end-state, and the ontology can grow in different directions beyond the control of system developers. Some common purposes of ontologies include the representation of knowledge in a given domain of interest, facilitating communication among system components, re-use by other applications, or as a common language for multiple systems within the same domain.

A good way for defining the purpose of the ontology is by means of use-cases. The intended use of the ontology can be broken down into specific, well-defined use-cases, in which actors and actions are identified, as well as the perceived components that will be involved in each action. Another tool for identifying the purpose of an ontology is a set of questions, which the ontology is supposed to answer.

The ontology is complete, in the context of a given set of requirements, when all the use-cases are supported by ontology concepts and all the questions to be asked can be answered by the current data, populating the ontology.

Step (2) - Building the ontology: Once the domain knowledge has been captured in free form, it is time to start building the ontology. This is the step, where the captured knowledge is formalized and concepts are given specific descriptive names to allow the communication with other stakeholders. The process of building the ontology can be described in the following steps:

1. Capture the knowledge in the domain of interest. Many knowledge acquisition techniques can be applied in this step, including textbooks, interviewing subject matter experts, databases of case studies, analysis reports, and so on. One of the primary methods of capturing knowledge in this domain is utilizing a subject matter expert to formalize the concepts and produce the model. Other sources of knowledge assist the expert in this task, such as books, military manuals, past plan analyses, and training material.
2. Identify the key concepts and relationships in the domain of interest. The key concepts are the ones that relate to the identified purpose of the ontology. They typically answer critical questions or contribute to the communication among system components and concepts that are involved in actions of use-cases. Other concepts that help to relate key concepts to each other or add details to key concepts, are considered supporting concepts. For example, the

- key concepts in a human factors ontology are likely to be: Person, Organization, Communication, Personal Traits, and Behavioral Traits.
3. Produce precise textual definitions of such concepts and relationships. The textual definitions help disambiguate the concepts and define their role in the ontology. Existing textual definitions in standard lexicons can help in this step. For example, the WordNet database offers a good electronic resource for common definitions of English language terms. The use of a lexicon like WordNet also facilitates the search for terms and their synonyms, for the purpose of analyzing free text. Other specialized lexicons, such as military manuals, can also be a good source for accepted definitions for common terms.
 4. Identify terms to refer to such concepts and relationships. The selection of good ontology terms helps developers understand the role of each concept and possibly the common uses of it. Also system developers do not need to go back to the formal definition of each term every time they need to use it. The selected terms should be expressive of the concept and close to its natural language description.
 5. Obtain agreement on all of the above. It is important for all stakeholders to agree on the selection of concepts and the terms used to refer to them in the ontology. Ontology-based systems are typically a collaborative effort, often among multiple organizations. To facilitate communication among all participants there has to be agreement on the ontology.
 6. Select a capture method (e.g., Protégé, UML, etc.). Modeling of ontologies is a step to formalizing the captured knowledge and producing an artifact, which can communicate that knowledge to other stakeholders. Most modeling methods have a graphical notation to easily connect concepts and navigate through the ontology. The criteria for selection of a modeling method are:
 - *Coverage*: Does the modeling method provide enough elements to represent all of the captured concepts and the types of relationships that exist among them?
 - *Granularity*: How much detail can the modeler represent in a concept?
 - *Learning curve*: Is this modeling method a standard method, which modelers are already familiar with? Or is it a new method that requires investment of time and effort to learn to use efficiently?

Protégé is the modeler of choice for OWL-based ontologies. There are other tools that support OWL development, such as Concept Maps, but the support that is offered by Protégé is stronger in visualization and ontology navigation.

Step (3) - Coding the ontology: The implementation of ontology-based systems requires translating the ontology model into an implementation language. The language chosen for

coding an ontology (e.g., formal logic, UML, OWL, etc.) has to provide the following characteristics:

- **Conceptual distance:** The ability of the language to represent abstract concepts at multiple levels of abstraction.
- **Expressive power:** The ability to represent complex concepts with consistent language constructs.
- **Standards compliant:** The language should follow accepted standards and notations to allow for better communication among development team members.
- **Translatability:** The language constructs have formal structures that can be converted to forms in other languages, without ambiguity.
- **Guidelines:** The model development process is supported by a set of guidelines and best practices.
- **Formal semantics:** The intended meaning of each language construct is unambiguous and well-defined.
- **Flexibility:** The ability to represent concepts in different ways, using different constructs.
- **User base:** The availability of user groups provides support for ontology development, through the exchange of experiences and best practices.
- **Availability:** The language has to be available, preferably, in the public domain, along with tools to support its use.

For example, the selected coding language may be OWL to facilitate communication with other system developers, especially in the case of a multi-organization effort. OWL satisfies many of the selection criteria mentioned above.

- **Conceptual distance:** OWL allows the representation of abstract concepts, maintaining its level of abstraction and allowing for details as needed.
- **Expressive power:** OWL employs description logic in a dynamic environment utilizing the open world assumption. Description logic is a powerful mechanism for stating concepts.
- **Standards:** OWL is based on RDF, which is a standard that is becoming more popular with many tools for processing formats.
- **Translatability:** As a formal language with well-defined semantics, OWL can be translated into other implementation languages, especially RDF-based languages. The degree to which the translation preserves all of the ontology features depends on the target language and its supported features.
- **Formal semantics:** OWL has well-defined semantics for language constructs. The semantics capability is supported by *Reasoner* specifications that describe what a valid structure should be.
- **Flexibility:** OWL offers a wide range of constructs to model concepts and relationships. In most cases, the modeler provides multiple choices to model

any concept. The selection of a particular construct is usually determined by the use-cases for the concept and the relationships to other concepts.

- *User base:* OWL enjoys a strong user base for OWL, especially using the modeler Protégé. There are also many conferences and user group meetings and on-line forums supporting the development of ontologies in OWL.
- *Availability:* The OWL specification is published in the public domain and tools for modeling in OWL are available for free (e.g., Protégé and CMAPtools).

The next step is to translate the model into actual system implementation. Two aspects need to be addressed in this step, namely: verification tools; and, code generation. Checking tools are needed to make sure that the ontology structure is consistent and remains consistent during system operation, after changes have been made. Code generation tools assist in taking a formal ontology consistently and repeatedly from a formal language to an implementation language. System implementation typically goes through multiple iterations that may require re-writing the basic model, or large sections of it. Utilizing code generation tools makes this task easier.

Step (4) - Integrating with existing ontologies: It is often the case that an ontology is being developed as an extension of an existing ontology or to connect with an existing ontology. In such cases, integration with the existing ontology must be carefully considered.

- Existence of other ontologies that are relevant to this ontology.
- All assumptions have to be made explicit.
- Agreement has to be achieved regarding concepts and relationships.

Step (5) - Evaluation: The ontology must be examined from a technical perspective, along with the associated software environment, and the documentation with respect to a frame of reference, which includes:

- Requirements specifications.
- Competency questions.
- Real world appropriateness.

The selection of the frame of reference and the evaluation criteria have to align with the purpose and requirements of the ontology. The semantic correctness of an ontology is crucial for the proper functioning of applications. In order to evaluate an ontology it is useful to employ a methodology that has two main components: structural analysis; and, domain knowledge analysis.

Structural Analysis: This involves the analysis of the structure of concepts in terms of hierarchy (taxonomy) and in terms of the relationships among concepts. The main criteria for this analysis are:

- *Uniqueness of concepts (no redundancy):* Every relevant concept in the domain should be represented in a clear and concise manner within the model. Concepts that are similar or have some common properties with other concepts should be represented in relationship to the existing concepts, either in a class hierarchy or through other types of relationships such as “part-of”.

The ease with which existing concepts co-exist and new ones can be added is a key indicator of the model's elegance and sophistication.

- *No circular reference should exist at any level:* Circular references can occur when a parent class in a class hierarchy inherits from a child class at any level down the hierarchy. This circular reference may not be obvious if the child class is more than two levels down from the parent class. Circular references are problematic because they confuse the semantics of the two concepts (e.g., "... a jet plane is a kind-of aircraft" and "... an aircraft is a kind-of jet plane").
- *Levels of Abstraction:* Class hierarchies can have any number of levels, where every level introduces more details to the classes at that level. The choice to add many attributes to a class in one level of the hierarchy or to create many levels with few attributes at each level has implications on the semantics of the model and on the operational aspects of applications that use this model.
- *Complexity (number of concepts + number of relationships for each concept):* The complexity of an ontology plays an important role in its usability. Applications typically traverse a collection of related concepts to form a context for reasoning or decision making. The more complex the ontology, the more involved it becomes for the application (and for the application developer) to form the proper context.

Domain Knowledge Analysis: Focuses on the *purpose* of the ontology. Use-cases for the application identify its information needs and form the basis for assessing the ontology's completeness. The criteria for this analysis are:

- *Coverage of use-cases (completeness):* Application use-cases define the different ways the ontology will be used. All concepts that are referenced in the target use-cases must exist in the ontology in some form (either directly or inferred). Other concepts not explicitly mentioned in any use-case may exist in the ontology serving as extended specifications for further reasoning or increased scope.
- *Partitioning:* The arrangement of classes in a hierarchy, where features of subclasses do not overlap, forms a *disjoint decomposition* of classes. When subclasses represent all the possible classifications of a super class, then this is called *exhaustive decomposition*. In this case, any instance of the super class is also an instance of one of the subclasses. These two properties of ontology partitioning (i.e., disjoint decomposition and exhaustive decomposition) place integrity constraints on the ontology and provide for tighter semantics, as well as a more powerfully expressive ontology that in turn leads to more straightforward reasoning capabilities.
- *Extensibility:* When the incorporation of additional concepts is required, perhaps due to the need to support additional use-cases, it should be possible to add these concepts without the need to re-structure the entire ontology. If engineered correctly, the incorporation of extended or entirely new concepts

can be achieved in a fairly isolated manner without unduly impacting unrelated areas of the model or actual model users.

- *Documentation:* The intended meaning and the usage of each concept must be clearly documented, so that reasoning facilities can effectively and appropriately employ them.

Step (6) - Documentation: The development of software components that are based on an ontology relies on good documentation of the ontology and the availability of the documentation to all developers. The documentation must include:

- Purpose and intended use of the ontology.
- Assumptions made at every level about concepts and their relationships.
- Primitives used to express the definitions (i.e., meta-ontology).
- Relationship to existing ontologies.

Using a lexicon such as WordNet standardizes the definitions across multiple developer teams and across organizations, and reduces the chances for ambiguity in dealing with concepts that may have multiple word-senses. The choice of Word-Net also offers the opportunity for other ontologies to integrate with the ontology under consideration, by examining the standard definition of its concepts and deciding on concept compatibility.

References

Barber K., A. Goel, D. Han, J. Kim, D. Lam, T. Liu, M. MacMahon, C. Martin and R. McKay (2003); 'Infrastructure for Design, Deployment and *Experimentation* of Distributed Agent-based Systems: The Requirements'; The Technologies, and an Example, Autonomous Agents and Multi-Agent Systems. Volume 7, No. 1-2 (pp 49-69).

Brown P. (2008); 'Implementing SOA: Total Architecture in Practice'; Addison-Wesley.

Erl T. (2008); 'SOA: Principles of Service Design'; Prentice Hall.

Fahad M. and M. Abdul Qadir (2008); 'A Framework for Ontology Evaluation'; Proceedings International Conference on Conceptual Structures (ICCS), Toulouse, France, July 7-11 (pp. 149-158).

Forgy C. (1982); 'Rete: A Fast Algorithm for the Many Pattern/Many Object Pattern Match Problem'; Artificial Intelligence, 19 (pp. 17-37).

Jennings N., K. Sycara and M. Wooldridge (1998); 'A Roadmap of Agent Research and Development'; Autonomous Agents and Multi-Agent Systems, Vol. 1 (pp. 7-38).

Wooldridge M. and N. Jennings (1995); 'Intelligent Agents: Theory and Practice'; The Knowledge Engineering Review, Vol. 10(2) (pp. 115-152).

Lu Q. and V. Haarslev (2006); 'OntoKBEval: A Support Tool for DL-based Evaluation of OWL Ontologies'; Proceedings OWL: Experiences and Directions 2006, Athens, Georgia, November 10-11.

Supekar K. (2005); 'A Peer-Review Approach for Ontology Evaluation'; Proceedings 8th International Protégé Conference, Madrid, Spain, July 18-21.

Wooldridge M., N. Jennings and D. Kinny (1999); 'A Methodology for Agent-Oriented Analysis and Design'; Proceedings Third International Conference on Autonomous Agents (Agents-99), Seattle, Washington.