Ontology

A collation by paulquek

Adapted from Barry Smith's draft @

http://ontology.buffalo.edu/smith/articles/ontology_PIC.pdf

Philosophical Ontology

Ontology as a branch of philosophy is the science of what is, of the kinds and structures of objects, properties, events, processes and relations in every area of reality. 'Ontology' is often used by philosophers as a synonym of 'metaphysics' (a label meaning literally: 'what comes after the *Physics*'), a term used by early students of Aristotle to refer to what Aristotle himself called 'first philosophy'. Sometimes 'ontology' is used in a broader sense, to refer to the study of what *might* exist; 'metaphysics' is then used for the study of which of the various alternative possible ontologies is in fact true of reality. (Ingarden 1964) The term 'ontology' (or *ontologia*) was coined in 1613, independently, by two philosophers, Rudolf Göckel (Goclenius), in his *Lexicon philosophicum* and Jacob Lorhard (Lorhardus), in his *Theatrum philosophicum*. Its first occurrence in English as recorded by the OED appears in Bailey's dictionary of 1721, which defines ontology as 'an Account of being in the Abstract'.

Download PDF file

http://ontology.buffalo.edu/smith/articles/ontology_PIC.pdf

*****

Adapted from Stanford Univ's KST Project @

http://www-ksl.stanford.edu/kst/what-is-an-ontology.html

[KST : Knowledge Sharing Technology]

What is an Ontology?

*By **Tom Gruber** <http://tomgruber.org> <gruber@ksl.stanford.edu>*

Short answer:

An ontology is ***a specification of a conceptualization***.

The word **"ontology"** seems to generate a lot of controversy in discussions about AI. It has a long history in philosophy, in which it refers to the subject of existence. It is also often confused with epistemology, which is about knowledge and knowing.

In the context of knowledge sharing, I use the term **ontology** to mean a *specification of a conceptualization*. That is, an ontology is a description (like a formal specification of a program) of the concepts and relationships that can exist for an agent or a community of agents. This definition is consistent with the usage of ontology as set-of-concept-definitions, but more general. And it is certainly a different sense of the word than its use in philosophy.

What is important is what an ontology is *for*. My colleagues and I have been designing ontologies for the purpose of enabling knowledge sharing and reuse. In that context, an ontology is a specification used for making ontological commitments. The formal definition of ontological commitment is given below. For pragmetic reasons, we choose to write an ontology as a set of definitions of formal vocabulary. Although this isn't the only way to specify a conceptualization, it has some nice properties for knowledge sharing among AI software (e.g., semantics independent of reader and context). Practically, an ontological commitment is an agreement to use a vocabulary (i.e., ask queries and make assertions) in a way that is consistent (but not complete) with respect to the theory specified by an ontology. We build agents that commit to ontologies. We design ontologies so we can share knowledge with and among these agents.

This definition is given in the article:

> **T. R. Gruber**. A translation approach to portable ontologies. *Knowledge Acquisition*, 5(2):199-220, 1993. Available on line <http://tomgruber.org/writing/ontolingua-kaj-1993.htm>.

A more detailed description is given in

T. R. Gruber. Toward principles for the design of ontologies used for knowledge sharing. Presented at the Padua workshop on Formal Ontology, March 1993, later published in *International Journal of Human-Computer Studies*, Vol. 43, Issues 4-5, November 1995, pp. 907-928. Available online <http://tomgruber.org/writing/onto-design.htm>.

With an excerpt attached.

Ontologies as a specification mechanism

A body of formally represented knowledge is based on a *conceptualization*: the objects, concepts, and other entities that are assumed to exist in some area of interest and the relationships that hold among them (**Genesereth & Nilsson**, 1987) . A conceptualization is an abstract, simplified view of the world that we wish to represent for some purpose. Every knowledge base, knowledge-based system, or knowledge-level agent is committed to some conceptualization, explicitly or implicitly.

An **ontology** is an explicit specification of a conceptualization. The term is borrowed from philosophy, where an Ontology is a systematic account of Existence. For AI systems, what "exists" is that which can be represented. When the knowledge of a domain is represented in a declarative formalism, the set of objects that can be represented is called the universe of discourse. This set of objects, and the describable relationships among them, are reflected in the representational vocabulary with which a knowledge-based program represents knowledge. Thus, in the context of AI, we can describe the ontology of a program by defining a set of representational terms. In such an ontology, definitions associate the names of entities in the universe of discourse (e.g., classes, relations, functions, or other objects) with human-readable text describing what the names mean, and formal axioms that constrain the interpretation and well-formed use of these terms. Formally, an ontology is the statement of a logical theory.[1]

We use common ontologies to describe *ontological commitments* for a set of agents so that they can communicate about a domain of discourse without necessarily operating on a globally shared theory. We say that an agent **commits** to an ontology if its observable actions are consistent with the definitions in the ontology. The idea of ontological commitments is based on the Knowledge-Level perspective (**Newell**, 1982) . The Knowledge Level is a level of description of the knowledge of an agent that is independent of the symbol-level representation used internally by the agent. Knowledge is attributed to agents by observing their actions; an agent "knows" something if it acts *as if* it had the information and is acting rationally to achieve its goals. The "actions" of agents --- including knowledge base servers and knowledge-based systems --- can be seen through a tell and ask functional interface (**Levesque**, 1984) , where a client interacts with an agent by making logical assertions (tell), and posing queries (ask).

Pragmatically, a common ontology defines the vocabulary with which queries and assertions are exchanged among agents. Ontological commitments are agreements to use the shared vocabulary in a coherent and consistent manner. The agents sharing a vocabulary need not share a knowledge base; each knows things the other does not, and an agent that commits to an ontology is not required to answer all queries that can be formulated in the shared vocabulary.

In short, a commitment to a common ontology is a guarantee of consistency, but not completeness, with respect to queries and assertions using the vocabulary defined in the ontology.

Notes:

[1] Ontologies are often equated with taxonomic hierarchies of classes, but class definitions, and the subsumption relation, but ontologies need not be limited to these forms. Ontologies are also not limited to conservative definitions, that is, definitions in the traditional logic sense that only introduce terminology and do not add any knowledge about the world (Enderton, 1972) . To specify a conceptualization one needs to state axioms that do constrain the possible interpretations for the defined terms.

****

Adapted from Wikipedia @

http://en.wikipedia.org/wiki/Ontology and http://en.wikipedia.org/wiki/Ontology_(computer_science)

In philosophy, **ontology** is *the study of being or existence* and forms the basic subject matter of metaphysics.

It seeks to describe or posit the basic categories and relationships of being or existence to define entities and types of entities within its framework.

Ontology can be said to study conceptions of reality; and, for the sake of distinction, at least to the extent to which its counterpart, epistemology can be represented as being a search for answers to the questions "What do you know?" and "How do you know it?", ontology can be represented as *a search for an answer to the question "What are the knowable things?"*.

Some philosophers, notably of the **Platonic** school, contend that all nouns refer to entities. Other philosophers contend that some nouns do not name entities but provide a kind of shorthand way of referring to a collection (of either objects or events). In this latter view, mind, instead of referring to an entity, refers to a collection of mental events experienced by a person; society refers to a collection of persons with some shared interactions, and geometry refers to a collection of a specific kind of intellectual activity.

Any **ontology** must give an account of which words refer to entities, which do not, why, and what categories result. When one applies this process to nouns such as electrons, energy, contract, happiness, time, truth, causality, and God, ontology becomes fundamental to many branches of philosophy.

In both computer science and information science, an **ontology** is a ***data model*** that represents a set of concepts within a domain and the relationships between those concepts. It is used to reason about the objects within that domain.

Ontologies are used in artificial intelligence, the semantic web, software engineering, biomedical informatics and information architecture as a form of knowledge representation about the world or some part of it. Ontologies generally describe:

Individuals: the basic or "ground level" objects
Classes: sets, collections, or types of objects[1]
Attributes: properties, features, characteristics, or parameters that objects can have and share
Relations: ways that objects can be related to one another
Events: the changing of attributes or relations

*****

Adapted from John F. Sowa @

http://www.jfsowa.com/ontology/

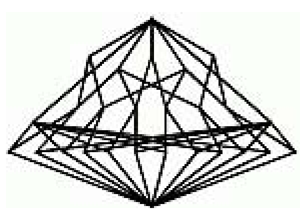

Words of Wisdom

There are more things in heaven and earth, Horatio,
Than are dreamt of in your philosophy.

William Shakespeare, *Hamlet*

The task of classifying all the words of language, or what's the same thing, all the ideas that seek expression, is the most stupendous of logical tasks. Anybody but the most accomplished logician must break down in it utterly; and even for the strongest man, it is the severest possible tax on the logical equipment and faculty.

Charles Sanders Peirce, letter to editor **B. E. Smith** of the *Century Dictionary*

The art of ranking things in genera and species is of no small importance and very much assists our judgment as well as our memory. You know how much it matters in botany, not to mention animals and other substances, or again moral and notional entities as some call them. Order largely depends on it, and many good authors write in such a way that their whole account could be divided and subdivided according to a procedure related to genera and species. This helps one not merely to retain things, but also to find them. And those who have laid out all sorts of notions under certain headings or categories have done something very useful.

Gottfried Wilhelm Leibniz, *New Essays on Human Understanding*

We must be systematic, but we should keep our systems open.

Alfred North Whitehead, *Modes of Thought*

Definition and Scope

The subject of ***ontology*** is the study of the ***categories*** of things that exist or may exist in some domain.

The product of such a study, called ***an ontology***, is a catalog of the types of things that are assumed to exist in a domain of interest ***D*** from the perspective of a person who uses a language ***L*** for the purpose of talking about ***D***.

The types in the ontology represent the *predicates*, *word senses*, or *concept and relation types* of the language ***L*** when used to discuss topics in the domain ***D***.

An uninterpreted logic, such as predicate calculus, conceptual graphs, or KIF, is *ontologically neutral*. It imposes no constraints on the subject matter or the way the subject may be characterized.

By itself, logic says nothing about anything, but the combination of logic with an ontology provides a language that can express relationships about the entities in the domain of interest.

1. An **informal ontology** may be specified by a catalog of types that are either undefined or defined only by statements in a natural language.

2. A **formal ontology** is specified by a collection of names for concept and relation types organized in a partial ordering by the type-subtype relation.

Formal ontologies are further distinguished by the way the subtypes are distinguished from their supertypes:

- an ***axiomatized ontology*** distinguishes subtypes by axioms and definitions stated in a formal language, such as logic or some computer-oriented notation that can be translated to logic;

- a ***prototype-based ontology*** distinguishes subtypes by a comparison with a typical member or *prototype* for each subtype.

Large ontologies often use a mixture of definitional methods: formal axioms and definitions are used for the terms in mathematics, physics, and engineering; and prototypes are used for plants, animals, and common household items.

KR Ontology

The ontology presented on this web site [**http://www.jfsowa.com/ontology/**] is based on the book ***Knowledge Representation*** by **John F. Sowa**.

The basic categories and distinctions have been derived from a variety of sources in logic, linguistics, philosophy, and artificial intelligence.

The two most important influences have been the philosophers **Charles Sanders Peirce** and **Alfred North Whitehead**, who were pioneers in ***symbolic logic***.

Peirce was also an associate editor of the *Century Dictionary*, for which he wrote, revised, or edited over 16,000 definitions. In calling that task "stupendous," he was looking beyond his personal experience of writing definitions in English to the task of stating complete definitions in logic, which he said was "a labor for generations of analysts, not for one." That labor, for which there was little practical application in the 19th century, is a major challenge for the 21st. *Without it, there is no hope of merging and integrating the ever expanding and multiplying databases and knowledge bases around the world.*

Yet as **Shakespeare** observed, *any philosophy is destined to be incomplete*. The continuing advance of science and human experience invevitably leads to new words and ideas that require extensions to any proposed system of categories. Whitehead's motto is the best guideline for any philosopher or scientist: "We must be systematic, but we should keep our systems open."

Hierarchies of Categories

To keep the system open-ended, the **KR ontology** is not based on a fixed hierarchy of categories, but on a framework of distinctions, from which the hierarchy is generated automatically.

For any particular application, the categories are not defined by drawing lines on a chart, but by selecting an appropriate set of distinctions. When the application-dependent distinctions are added to the basic set, a new **lattice** of categories can be created by pushing a button.

The icon below

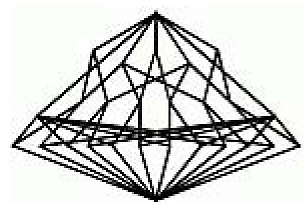

illustrates the lattice used to represent the ***top-level categories***, but lattices can also be used to represent categories at any level.

As an example of a lattice of lower-level types, **Figure 1** shows beverage types classified according to the attributes *alcoholic, nonalcoholic, hot, sparkling, caffeinic, madeFromGrapes,* and *madeFromGrain*. This lattice was derived from the attributes by the method of ***formal concept analysis***.

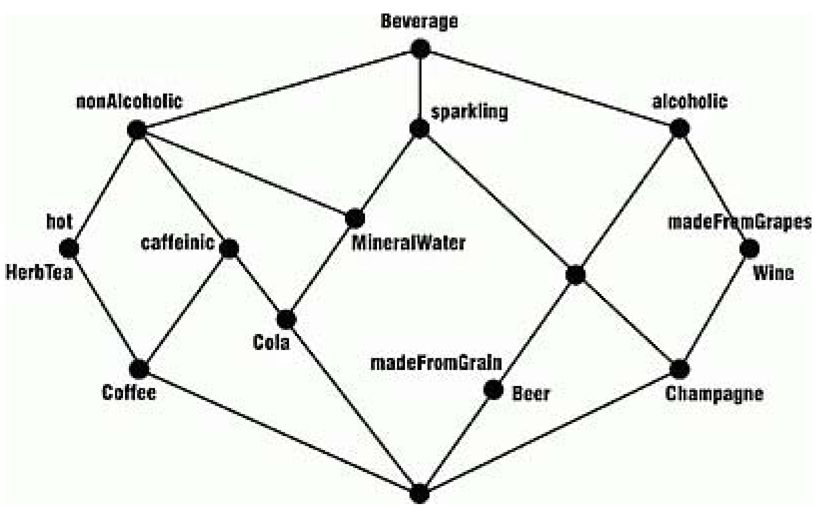

Figure 1: A lattice constructed by the method of formal concept analysis

The FCA techniques belong to the general class of *data mining* procedures, which find patterns in a relational database. The raw data used to generate FCA lattices is the same kind of data that could be used for other data mining techniques, such as neural networks. Each technique has its own advantages and disadvantages, depending on how the result is going to used. For ontology, the FCA technique produces a sublattice that can be automatically merged with a more general lattice of categories. In the case of **Figure 1**, the top node represents the type Beverage, which could be defined as DrinkableLiquid in terms of higher-level categories.