Possible-Worlds Ontology for the Semantic Web

Ahti Pietarinen
Department of Philosophy
University of Helsinki
P.O. Box 9, FIN-00014 University of Helsinki
pietarin@cc.helsinki.fi

Abstract

The emergence of the Semantic Web suggests new perspectives to knowledge acquisition and representation, increasing the usefulness of possible-worlds semantics of epistemic logic (knowledge, belief and information). In this poster, it is argued that possible worlds are a foundation of the Semantic Web, because one can (i) address the problem of the property of omniscience in logical approaches to knowledge, (ii) develop question-answering systems, interrogation and the semantics of questions, and (iii) see what the role model theory plays is in the Semantic Web research.

Keywords: Possible-worlds semantics, the Semantic Web, epistemic logic, knowledge, semantics of questions.

1 Knowledge representation and the Web

The primary goal of the Semantic Web vision [1, 2, 3] is to enable intelligent queries for knowledge on the Web instead of the conventional string matching.

To do this, we need to accomplish the following:

1. We need to understand the semantic mechanism of all kinds of questions, and what kind of components the process of questioning the Web formally consists of.
2. We need to rigorously capture, represent or symbolise the knowledge contained on the Web.

We propose to use epistemic logic in accomplishing these, together with the affiliated fields of knowledge acquisition, representation and interrogation.

The connection between the methods of inquiry and interrogation is knowledge acquisition, 'the Semantic Web mining'. The important issue is the semantics of questions.

In order to develop viable queries on the Web, we need to investigate the totality of the question-answer relationship. This brings the role of semantic information to the fore in the Web research.

2 Epistemic logic and knowledge representation

2.1 What is it?

Epistemic logic provides a theory that can be used for both knowledge representation and knowledge acquisition. It addresses the problems of logical omniscience, introspection and reflexive epistemic logic, representing metalevel knowledge, non-monotonic versions of autoepistemic logic, the interplay between knowledge and belief, and other epistemic notions including awareness and implicit versus explicit knowledge.

The trade-off between expressivity and performance is a long-standing issue whose import depends largely on the methodological viewpoints that one is adopting.

The basic idea is that knowledge is represented by a modal operator K, attached to a proposition φ and subscripted by an index denoting the agent. $K_i\varphi$ means that 'the agent i knows φ'. The central concept is that of a possible world, plus the relations between possible worlds.

The truth-clauses are relativised to the set of worlds $\mathcal{W}$. Let M be a model for $\mathcal{L}$-formulas φ, $w_0 \in \mathcal{W}$: $w' \in [w_0]_{\rho_i}$ means that w' is i-accessible from w_0 [4].

2.2 The problem of logical omniscience

There is a conceptual problem in the possible-worlds semantics: *logical omniscience*. Since the traditional possible worlds semantics treats knowledge as necessity ('if $\models \varphi$ then $\models K\varphi$') and since normal modal systems contain the axiom **K** ('$K(\varphi \to \psi) \to (K\varphi \to K\psi)$'), it immediately follows that every formula implied by a known formula is also known.

The property of logical omniscience thus says that whenever an agent knows the formula φ ($K\varphi$), and φ logically implies the formula ψ ($\varphi \models \psi$) wrt the class $\mathbf{Mod}(M)$ of models, then the agent knows ψ, that is, $K\psi$.

Logics extending **K** and interpreted by possible worlds are afflicted by this: agents' attitudes are ideal in that they end up knowing all the logical consequences, which is not consistent with rationality.

2.3 Variety of knowledge constructions

Because of this, we need to distinguish between different senses of knowledge.

First, knowledge that is contained in the Web (the Oracle) is passive, *implicit knowledge*. We represent it by a set of propositions Φ. We assume that the User has unrestricted access to this set in that all answerable questions that the User puts to the Oracle are found in Φ. The propositions in Φ are not known to the User, however, apart from those that are the outcomes of the User's process of questioning of the Oracle.

Secondly, *explicit knowledge* will be a set of propositions Ψ that the User knows as a result of the questions put to the Oracle and successful answers to them. A further restriction is between knowledge in this sense and explicit knowledge that is true, that is, has a model in the domain of discourse. Alternatively, explicit true knowledge can be viewed as the User's range of awareness over Ψ.

Full logical omniscience says that whenever an agent knows the formula φ ($K\varphi$), and φ logically implies the formula ψ ($\varphi \models \psi$) with respect to the class $\mathbf{Mod}(M)$ of models, then the agent knows ψ, that is, $K\psi$.

This problem is also a computational one: certainly an artificial agent cannot compute all logical consequences of its beliefs, for the agents are resource-bounded by computation time and memory space.

2.4 Avoiding omniscience

Logical omniscience is a property that has its roots in the very nature knowledge is defined using possible worlds, namely as truth in them. For this reason, some stronger variants are needed.

An *impossible worlds model* is a 4-tuple $\mathcal{M} = \langle \mathcal{W}, \mathcal{W}^*, \rho, V \rangle$, where $\mathcal{W}$ is a non-empty set of possible worlds, $\mathcal{W}^*$ is a set of impossible worlds, ρ is an accessibility relation on $(\mathcal{W} \cup \mathcal{W}^*) \times (\mathcal{W} \cup \mathcal{W}^*)$, and $V\colon (\Gamma \times (\mathcal{W} \cup \mathcal{W}^*)) \to \{0, 1\}$ is a valuation assigning truth-values to formulas Γ in all possible and impossible worlds.

Valuations are no longer recursively defined. However, an agent must still consider impossible worlds when evaluating formulas and hence, depending on the choice of such worlds and valuations, may fail to believe the logical consequences of knowledge.

Impossible worlds semantics solves logical omniscience for good: no matter what closure properties are put forward, the truth conditions for the connectives in impossible worlds can be chosen such that no closure is generated.

3 The Semantic Web and modal model theory

3.1 Models for what?

The Internet is a semiotic system with signs and symbols, and programs that manipulate them. The machine understanding of the Web content is a challenge.

But wherein do we find the medium that relates objects to their representations?

The proposal: Semantics can be found in its relation to the Web analogously to the way it is related to logic. That is, in order to build a Semantic Web we need a sensible *model theory* and relate our the preferred standardised language to any given model. Thus, the Web is our model, and languages such as XML, KIF, RDF(S) or DAML+OIL comprise our logics.

In particular, OIL is a language built on top of the existing W3C standards of RDF and XML, related to the description logic SHIQ. The idea of merging description logic and Web-languages is to create a core that takes care of ontological and epistemological needs. Ontologies are very much what the languages support, both syntactically (Web-based language) and semantically (description logics).

Epistemologies are something that are tried to be derived from knowledge representation schemes. For instance, in OIL the preferred epistemology is to use a frame-based system.

The model theory in [8] suggests to assign meanings to expressions of RDF and RDF(S). They are directed RDF graphs, nodes labelled with URIs or literals, and edges labelled with URIs. URIs at edges correspond to predicates or relations of logic, and at nodes to the names of individuals. Literals are a subset of constants syntactically distinguished from URIs. Unlabelled nodes correspond to variables, and labelled edges to predicates or relations between expressions.

This gives a model in which expressions of a given language are evaluated in the sense of model theory.

Possible worlds do some extra work here. Worlds are URIs, and accessibilities between different URIs are simply given by other URIs. This enriches the language of RDF(S) by modal expressions whose intended interpretation is that some URIs 'can be necessarily found' or 'be possibly found'.

An alternative interpretation is an epistemic one, capturing that an agent 'knows that a proposition φ holds' or dually that 'it is possible, for all that an agent knows, that a proposition φ holds'.

3.2 Consequences for the Semantic Web

We still do not quite know what a Semantic Web is. An attempt to define is undertaken by the Semantic Web Agreement Group [7]:

> The Semantic Web is a Web that includes documents, or portions of documents, describing explicit relationships between things and containing semantic information intended for automated processing by our machines.

But surely we do not mean to define the Semantic Web as one that contains semantic information. It has to be more: a model for expressions. It is a model for the data that is defined through shared ontologies and processed by an automated system.

The relations denote various accessibilities and documents or things are what comprise the 'possible worlds' (the sets of propositions or theories). To avoid omniscience (not a welcome feature of any Semantic Web), varieties of knowledge and impossible possible are taken into account.

3.3 Interrogative games

Users interrogate the Oracle. This questioning is an *extensive game* with a partitional structure of players' information. The reason for using games is that games codify both the position and the information available at any location.

This affects explicit knowledge, which agent knows as an outcome of the *interrogative process.* We assume that in some instances an agent is attuned to these outcomes. A position or a state in a game captures what the agent's current active knowledge is.

The notion of implicit knowledge is not locational in this sense, because it is just what the Oracle contains, not being structured and incorporated into the evolving notion of a game.

Yet implicit knowledge is known to the User in the sense that he or she can choose suitable elements from it. This is achieved by strategies that the players have an access to. Often the range is restricted, for example when there is only partial knowledge of previous answers, so that the User's choice is restricted to just some alternative.

There are two kinds of answers that the User can elicit from the oracle: deductive and interrogative ones. This comes close to the interrogative approach to inquiry [5].

3.4 Answers to questions

The enormous growth of the amount of data in the Web fosters interdisciplined, holistic approaches to the interplay between the Client and the Oracle.

This is what the interrogative approach addresses. In order to do this, one needs a theory of questions plus what counts as a *conclusive answer* to a question.

The semantics is that a 'wh'-subset (who, what, where...) of questions is an epistemic imperative, like 'Bring it about that I know wh-'. Hence an answer would be a statement of epistemic logic that can symbolise statements such as 'I know what'.

The Semantic Web that is capable of interpreting questions by epistemic sentences is needed in question-answering. This in turn again needs 'possible worlds'.

Also, in the Web *presuppositions* are implicit knowledge. A question cannot be answered unless the answer is in principle found among the set of propositions that constitute implicit knowledge, because this set determines what questions are possible in the first place.

3.5 Epistemic logic versus description logic

In *description logic*, the basic primitives are concepts denoting classes and entities, and roles denoting properties and relations between entities. Individuals are instances of classes, and intersection (conjunction) is used in creating complex concepts.

From an abstract point of view, descriptions logics agree with propositional modal logic, which in turn agrees with bounded fragments of first-order logics. Hence epistemic logic suffices — indeed, description logics have been used to go proxy for epistemic concepts [6].

As soon as predicate extensions of epistemic concepts are needed, such as in questioning the Web, description logics lose their expressive capacities, however.

4 Conclusion

A useful definition of the notion of information is that the more informative the sentence is the more alternatives it excludes: an informative sentence has a low probability. The probability of a statement is the proportion of worlds consistent with the knowledge in which the statement is true, and the information is increased with the number of worlds that are ruled out in learning that the statement is true.

The recent Web technology should take all aspects of semantic information into account. We already have an efficient medium for information transmission, but the advances in the semantic side of languages will increasingly put pressure on the 'possible states' theories of information.

References

[1] Berners-Lee, T., 1999. *Weaving the Web*, London: Orion Business Books.

[2] Decker, S. et al., Knowledge representation on the Web, in Baader, F. (ed.), *International Workshop on Description Logics (DL'00)*, `http://www.cs.vu.nl/~frankh/abstracts/DL00.html`.

[3] Fensel, D. and Musen, M.A., 2001. The Semantic Web: A brain for the mankind, *IEEE Intelligent Systems* March/April 2001, 24–25.

[4] Hintikka, J., 1962. *Knowledge and Belief*, Ithaca: Cornell University Press.

[5] Hintikka, J., 1989. Knowledge representation and the interrogative model of inquiry, in Clay, M. and Lehrer, K. (eds), *Knowledge and Skepticism*, Boulder: Westview Press, 155–183.

[6] Hustadt, U., 1995. Introducing epistemic operators into a description logic, in Laux, A. and Wansing, H. (eds) *Knowledge and Belief in Philosophy and AI*, Berlin: Akademie Verlag, 65–86.

[7] Semantic Web Agreement Group, 2001. What is the Semantic Web?, `http://swag.webns.net/whatIsSW`.

[8] W3C Consortium, 2001. RDF model theory, `http://www.w3.org/TR/2001/WD-rdf-mt-20010925/`.