

Logical Systems: Towards Protocols for Web-Based Meaning Negotiation

James Farrugia

National Center for Geographic Information and Analysis
Department of Spatial Information Science and Engineering
University of Maine, Orono, ME 04469-5711, USA
jim@spatial.maine.edu

Abstract

Our thesis is that before Web-based agents can negotiate meanings, they need to agree on high-level protocols based on *logical systems*. The Semantic Web community is setting the stage for semantic interoperability among Web-based software agents by developing standard languages with well-defined semantics. But exactly how these languages might be used to negotiate meanings is an open question. The agent-oriented perspective of negotiation as a kind of deal-making does not seem to mesh with the perspective of Semantic Web community, which seems more inclined to fix meanings, rather than bargain over them. Fundamental to treatments of meanings on the Semantic Web is the ability of agents to perform logical inferences; therefore, any approach to Web-based meaning negotiation should, take into account those aspects of logic that have a bearing on logical inferences. *Logical systems* provide a general way of accounting for those aspects of an agent's *logical setup* that have a bearing on inferences. An explicit taking into account of agents' logical systems provides a suitably rich foundation on which to base protocols for meaning negotiation on the Web.

Introduction

This workshop explores an approach to semantic interoperability called *meaning negotiation*, in which autonomous agents negotiate agreements about the content and purpose of their interaction, rather than refer to some shared conceptualization of a domain for resolution of semantic and pragmatic issues. But before agents can negotiate anything about semantics or pragmatics, they need to agree on the protocols that govern their interactions. Negotiation protocols are high-level protocols that deal with the content of agent interactions, as well as with rules that describe "the sequence of offers and counter-offers that are allowed" (Rosenschein and Zlotkin 1994). In the case of meaning negotiation, these protocols govern both the meanings to be negotiated and the rules by which this negotiation is carried out.

Rules aside, the idea of agents actually *negotiating* meanings takes some getting used to. The reason is that a significant intuition runs up against an important tradition. From the agent-oriented perspective, the term *meaning negotiation* suggests that meanings are the contents over which a bargain of some sort can be struck. Seen from this perspective, meaning *negotiation* is a utility-based process that ends in success when agents work out a mutually acceptable deal (Müller 1996). This view reflects the

intuition that meanings depend on an interaction between agents that takes into account bargaining, context, and pragmatic considerations. But this intuition seems to run counter to the model-theoretic tradition in logic, which aims to *fix* meanings through well-defined mappings (Genesereth and Nilsson 1987; Manzano 1999). Thus, there is a tension between the agent-oriented intuition that treats meanings as the contents about which some pragmatic deal may be struck, and the logical tradition that deals with meanings as well-defined mappings.

It is fair to say that current treatments of semantics on the Web fall within the logical tradition. The *Semantic Web* is "a vision for the future of the Web in which information is given explicit meaning, making it easier for machines to automatically process and integrate information available on the Web" (Heflin *et al.* 2002). In this vision, "even agents that were not expressly designed to work together can transfer data among themselves when the data come with semantics" (Berners-Lee *et al.* 2001). Although the success of this vision depends on semantic interoperability among agents (i.e., the ability of agents to work together when they may be using different conceptualizations of their domains), the idea of meaning negotiation as a possible approach to semantic interoperability has apparently not taken hold in the Semantic Web community. The reason is not hard to see. A central tenet of Semantic Web developers is that languages used to carry meanings on the Web should "have a well defined formal semantics" (Horrocks 2002), which does not permit ambiguity and which allows for logical inferences. One principal way of providing Web-based languages with a formal semantics is via model-theoretic semantics (Patel-Schneider 2002). Thus, developers of Semantic Web languages are aligning themselves with a logical tradition, which seeks to fix meanings via well-defined correspondences.

At the same time, the Semantic Web community recognizes that different agents may employ different meanings in the conceptualization of their domains. Thus, it is clear that different sets of meanings will need to be accommodated on the Web. But will different sets of meanings be *negotiated*? That is, can the logical-semantic tradition of well-defined semantics for the Web accommodate the agent-oriented perspective of meaning negotiation? Yes, but not before agents agree on the protocols that govern their interaction. In this abstract we focus on the logical foundations for these protocols. We leave to other researchers in this workshop (and to our own

future work) the specific treatment of *how* agents might adopt a meaning-negotiation approach to semantic interoperability on the Web.

Our claim is that a recognition of each agent's *logical system* (Meseguer 1989) forms the proper basis for meaning-negotiation protocols on the Web. A logical system is an axiomatic treatment of the following logical components: sentences in a formal language; consequence relations over these sentences; model theory (and its generalization to institutions); and a proof calculus. We show that the kinds of meanings relevant to agents on the Web include not just the model-theoretic components of a logical system, but also other components that have a bearing on inferences. We recommend, therefore, that protocols for Web-based meaning negotiation take into account all relevant components of agents' logical systems. The actual negotiation of meanings, we expect, will be a back-and-forth process that ends in mutual recognition of what kinds of mappings are possible between agents' logical systems.

In the next three sections we explain our understanding of *meaning negotiation*, *semantics on the Web*, and *logical systems*. We follow these sections with a discussion of how the components of a logical system are relevant to the treatment of inferences on the Web. We conclude by reviewing the main points and discussing some loose ends.

Meaning Negotiation

Software agents interact in public and private ways according to particular rules. The public behaviors of agents are governed by negotiation protocols. These protocols "specify the possibilities of initiating a negotiation cycle and of responding to a message" (Muller 1996). The private behavior of an agent consists of strategies that direct it to select some possible deal(s). These strategies are usually based on optimizing some utility function that governs the agent's choice of a deal (Rosenschein and Zlotkin 1994). The "space of possible deals" (Rosenschein and Zlotkin 1994), i.e., the content of the communication, is the set of options over which agents negotiate.

For example, consider two agents negotiating about the selection of one or more restaurants from a set of restaurants. Factors influencing the agents' utility functions might include price, type of fare available, distance from present location, etc. Agent 1 might have a simple utility function that directs the choice of restaurant based on the minimization of total price. Agent 2's utility function might direct a choice based on the broadest range of available cuisine. If these agents are people, then their public behavior (as shown through their interaction protocol) might simply be conventionally accepted norms of human discourse, which typically will not even bear mention or notice. The content of the negotiation is the set of possible restaurants from which one restaurant may be selected. The private behavior would then be the strategy

of optimizing one's utility function. In this simple example, a deal can be struck (a choice of restaurant can be made) if agents can agree on one or more selections that optimize their individual utility functions.

Contrast this scenario with one in which software agents are negotiating about meanings rather than about restaurants. Suppose the agents use the same formal language (with the same model-theoretic semantics) to represent their meanings. In this case, what meanings exist to be negotiated? Not the meanings fixed by the model theory. Rather, the kinds of meanings open to negotiation will be those meanings that reflect each agent's idiosyncratic view of the world. Such a world view is typically encoded by axioms that specify certain properties of, or relationships among, entities in a particular domain. Thus, we can envision a scenario in which agents use a common language and model theory to express different sets of axioms that describe their individual conceptualizations of the world.

Semantics on the Web

The Semantic Web is moving towards such a scenario, though it is not yet clear exactly how agents might work together when they are operating with different conceptualizations of a given domain. The move towards a common language is reflected through the development of standard languages such as OIL (Horrocks *et al.* 2000), DAML+OIL (van Harmelen *et al.* 2001), and the latest Web Ontology language, still under development, called OWL (Heflin *et al.* 2002). Each of these languages has a formal semantics. Any of these languages, together with its semantics, constitutes a tool that can be used to write meaningful statements for a domain of interest for which the language constructs are applicable. For instance, once OIL's language constructs (class definitions and slot definitions) are translated to the *SHIQ* description logic, OIL's model theory provides a way to give meaning to this translated terminology. The "meaning of a *SHIQ* terminology, and of the common inference problems, is given in terms of a Tarski style model theoretic semantics using *interpretations*" (Horrocks *et al.* 2000). Note that the above quotation reveals two important aspects of meanings that are significant for the treatment of meanings on the Web. First, there are denotational mappings (interpretations of terminology). Second, by means of these mappings and the definition of a satisfaction relation, there are a language's inferential capabilities. The attitude towards meanings by the Semantic Web community might well be summarized: "if a language does not have a well-defined semantics, it is of limited use, since machines will not be able to do automated reasoning with it."

The focus on a single language still leaves open the question of how different agents, each using the same language, can interoperate when they have different conceptualizations about one or more domains. Also left open is the question of semantic interoperability when the agents use different languages to represent their

conceptualizations. Further, since we have seen that meanings on the Web deal with logical inferences, one needs to consider how one agent's valid inferences relate to another's. Instead of directly attacking the first question, we indirectly attack them all, by framing the debate in terms of the *logical setup* that agents must agree about before they can deal with any kind of semantic interoperability. We now turn to the notion of logical system, which formalizes the intuitive notion of an agent's *logical setup*.

Logical Systems

We give a cursory and succinct description of logical systems, omitting the technical details, which can be found in Meseguer (1989) and Martí-Oliet and Meseguer (1994). The goal of our presentation is to make plausible the arguments that follow in the next section. A logical system is a modular theory of general logics that treats the following notions: syntax; entailment of a sentence from a set of sentences; model and satisfaction of a sentence by a model; and proof calculus.

The **syntax** of a logic is given by a signature (consisting, for example, of relation and function symbols), a logical vocabulary, and a grammar, which taken together prescribe how all *sentences* of the logic can be constructed.

An **entailment system** axiomatizes the consequence relation of a logic, which holds between a set of sentences that are assumed (e.g., axioms) and a sentence that is proved (the conclusion). Sometimes this consequence relation is called the *provability* relation. (Note also that some authors use the word *entailment* to refer the *semantic* consequence relation, not the *syntactic* consequence relation.)

The notions of **model** and **satisfaction** are covered by institutions (Goguen and Burstall 1984; Goguen and Burstall 1992). An institution "consists of a category of signatures such that associated with each signature are sentences, models, and a relationship of satisfaction that, in a certain sense, is invariant under change of signature" (Goguen and Burstall 1992). Whereas "traditional model theory assumes a fixed vocabulary, institutions allow us to consider many different vocabularies at once" (Goguen and Burstall 1992).

A **logic** "is then obtained by combining an entailment system with an institution in such a way that a soundness condition (relating provability and satisfaction) holds" (Meseguer 1989).

Finally, by choosing a particular **proof calculus** to go along with a given logic, one obtains a **logical system** (Meseguer 1989).

Towards Protocols for Web-Based Meaning Negotiation

A principal reason that the Semantic Web community treats meanings via well-defined semantics is to enable machines to carry out logical inferences. Thus, if some

aspect of logic affects logical inferences, this aspect is at least potentially relevant to the treatment of meanings on the Web. In particular, we claim, if two agents attempt to negotiate meanings on the Web, then the aspects of each agent's logical system that concern logical inferences are relevant to the process of meaning negotiation.

Consider two agents with different entailment systems. These entailment systems will not necessarily axiomatize the consequence relation in the same way, though many axiomatizations will stipulate that the consequence relation satisfy the properties of reflexivity, monotonicity, and cut.

A consequence relation is *reflexive* if and only if when the conclusion is one of the assumptions, then the conclusion follows from the assumptions; it is *monotonic* if and only if when the conclusion follows from a set of assumption, that conclusion also follows from those assumptions augmented with additional assumptions; it has the *cut* property if and only if when a conclusion from a set of assumptions is added to this set, and this larger set is used to make a second conclusion, then this second conclusion can be reached from the original assumptions (Gabbay 1998).

Some logical systems may stipulate less of their consequence relation (e.g., that monotonicity does not always hold) while others may stipulate more (Gabbay 1998). The axiomatization of the consequence relation affects inferences, and thus it has a bearing on how agents negotiate meanings on the Web. We recognize that many developers in the Semantic Web community would argue against allowing agents to use nonmonotonic reasoning. We include this example to illustrate the range of possibilities that agents might use, and thus what needs to be accounted for in the logical foundations of negotiation protocols.

Consider now the case of agents using different signatures. To interoperate, they will likely need a way to relate their models. The "essence of the institution notion is that a change of signature (by a signature morphism) induces *consistent* changes in sentences and models" (Goguen and Burstall 1992). Both Meseguer (1989) and Goguen and Burstall (1992) provide techniques to map between different institutions. Differences between institutions clearly deal with meanings; thus, they also clearly impact the kinds of valid inferencing that agents can perform, either individually or jointly.

Finally, consider agents that use different proof calculi, say, for the same underlying logic. Although the entailment relation is "precisely what remains *invariant* under [different] proof calculi" (Meseguer 1989), and thus it would seem that the proof calculus itself has little bearing on inferences, it is still quite relevant from a practical point of view.

All three major components of logical systems may, in any given situation, need to be dealt with in order to allow agents to interoperate.

We believe that in the future agents will likely use whatever language, model theory, and proof calculus that

suit their needs. Having a general way to handle this variety of *logical setups* will enable the adequate treatment of semantic interoperability on the Web. Logical systems are an appropriate foundation on which to base protocols for meaning negotiation on the Web, and mappings between logical systems will be key in developing these protocols.

Summary

Meanings on the Semantic Web are currently being dealt with through standard languages, each equipped with its own model-theoretic semantics. The model-theoretic treatment of semantics fits within a logical tradition that deals with meanings through well-defined mappings. Of key importance to the Semantic Web community is a treatment of meanings that enables automatic reasoning by machines. This community appears to be advocating the position that although agents may have different axiomatizations of their domains, they will likely use the same *lingua franca* to formulate these axiomatizations.

But just as we cannot suppose that agents will use the same ontologies to understand the world, we cannot reasonably suppose that agents will use the same logical system in dealing with their worlds. Fundamentally, meaning negotiation on the Web needs to take into account all aspects of logic that have a bearing on inferences. Once agents agree on a suitably general view of logic, they will then be in a position to negotiate meanings, perhaps via mappings between their own individual *logical setups*.

Meseguer's logical systems are one framework that formalizes the aspects of logic that have a bearing on meanings, providing a suitable foundation on which to construct protocols for meaning negotiation on the Web. So, agents need to *begin* the process of meaning negotiation by first recognizing each other's logical system, in order to determine what kinds of meaningful interactions are possible. The negotiation itself will likely then be based on possible mappings between the agents' logical systems.

We recognize that the sketch presented in this abstract leaves many loose ends, some of which we list below. We hope to address them in future work.

- other views of logical systems should be addressed (Gabbay 1994), and a justification for adopting Meseguer's approach should be given;
- Meseguer's later work on formal interoperability should be considered (Meseguer 1998);
- detailed examples should be put forth of a possible protocol for meaning negotiation and how this negotiation could then be based on mappings between logical systems;
- we completely sidestepped issues of context and the effects of pragmatic considerations on meaning negotiation;
- we need to integrate our views with related work on agent communication languages.

Acknowledgments

This work was partially supported by the National Imagery and Mapping Agency under grant number NMA202-97-1-1023 (PI: Max Egenhofer).

References

- Berners-Lee, T., Hendler, J. and Lassila, O. 2001. The Semantic Web. *Scientific American* 184(5): 34-43.
- Gabbay, D. 1998. *Elementary Logics: A Procedural Perspective*. London: Prentice Hall Europe.
- Gabbay, D. ed. 1994. *What is a Logical System?* Oxford: Oxford University Press.
- Genesereth, M. and Nilsson, N. 1987. *Logical Foundations of Artificial Intelligence*. San Francisco: Morgan Kaufmann,
- Goguen, J. and Burstall, R. 1992. Institutions: Abstract Model Theory for Specification and Programming. *Journal of the ACM* 39(1): 95-146.
- Goguen, J. and Burstall, R. (1984) Introducing Institutions. In E. Clarke and D. Kozen, eds. *Logics of Programming Workshop*. 221-256. Berlin: Springer.
- Heflin, J., Volz, R. and Dale, J. eds. 2002. Requirements for a Web Ontology Language. <http://www.w3.org/TR/webontreq/>
- Horrocks, I. 2002. An Ontology Language for the Semantic Web. *IEEE Intelligent Systems*, (forthcoming).
- Horrocks, I., Fensel, D., Broekstra, J., Decker, S., Erdmann, M., Goble, C., van Harmelen, F., Klein, M., Staab, S., Studer, R., and Motta, E. 2000. The Ontology Inference Layer OIL. <http://www.cs.vu.nl/~dieter/oil/Tr/oil.pdf>
- Manzano, M. 1999. *Model Theory*. Oxford: Oxford University Press.
- Martí-Oliet, N. and Meseguer, J. 1994. General Logics and Logical Frameworks. In D. Gabbay. ed. *What is a Logical System?* 355-391, Oxford: Oxford University Press.
- Meseguer, J. 1989. General Logics. In H.-D. Ebbinghaus, J. Fernández-Prida, M. Garrido, D. Lascar, and M. Rodríguez Artalejo. eds. *Proceedings of Logic Colloquium '87*, 275-329, Amsterdam: North Holland.
- Meseguer, J. 1998. Formal Interoperability. In The 1998 Conference on Mathematics in Artificial Intelligence, Fort Lauderdale, Florida.
- Müller, H. 1996. Negotiation Principles. In G. O'Hare and N. Jennings, eds. *Foundations of Distributed Artificial Intelligence*, 211-229. New York: Wiley.
- Patel-Schneider, P. 2002. Building the Ontology Layer of the Web. *IEEE Intelligent Systems* (forthcoming).
- Rosenschein, J. and Zlotkin, G. 1994. *Rules of Encounter: Designing Conventions for Automated Negotiation among Computers*. Cambridge: MIT Press.
- van Harmelen, F., Patel-Schneider, P. and Horrocks, I. 2001. Reference description of the DAML+OIL (March 2001) ontology markup language. <http://www.daml.org/2001/03/reference.html>