

Meaning Negotiation: an invitation

Paolo Bouquet

Dept. of Information and Communication Tech.
University of Trento (Trento, Italy)
bouquet@ict.unitn.it

Massimo Warglien

Dipartimento di Economia e Direzione Aziendale
University of Venice (Venice, Italy)
warglien@unive.it

The development of distributed applications over large networks of computers raises the issue of semantic heterogeneity between autonomously developed sources of information or service providers. If, on the one hand, it is generally recognized that we need a semantic layer to provide more and more complex services, on the other hand it is clear that such a layer requires the ability to “bridge” (interoperate) representations whose heterogeneity is much deeper than that of syntactic formats. As a matter of fact, different people (individuals, groups, communities, service providers, corporations, and so on) view – and thus represent – the world from different perspectives, using different interpretation schemas, having different goals in mind. This diversity of viewpoints is the real source of semantic heterogeneity; it is intrinsic of the way cognitive agents (in particular, human beings) know the world, and as such it cannot be wiped away.

Semantic heterogeneity manifests itself in different forms. Simple examples are:

Synonymy and polysemy. The simplest source of semantic heterogeneity is the fact that people use different words to refer to the same thing (*synonymy*), or the same word to mean different things (*polysemy*). Synonymy and polysemy may seem simple problems, as we only need to check a good dictionary to find that two different words have the same meaning, or that a word may have more senses. However, we can extend the notions of synonymy and polysemy to much more complex situations, in which linguistic meaning is not sufficient to solve the problem. For example, some people claim that “holonic systems” and “agent-based systems” refer to the same kind of system in different domains; or that “knowledge management” is a more fashionable name of what a few years ago, perhaps in different academic circles, was called “organizational learning”. The interesting thing is that people might disagree on the truth of these claims, this way showing that a broader notion of synonymy may require some degree of world knowledge.

Partiality and granularity. Imagine we ask a travel agent, an hotel manager, and a tourist to conceptualize a simple service, say booking a room. It goes without saying that

the descriptions will be very different, even if the service is the same. Each of them will stress different aspects, will consider notions which are completely disregarded by the others, will describe in great detail things that the others consider almost irrelevant. Consequently, trying to map their conceptualizations onto each other can be very hard. Unlike the previous form of heterogeneity, not only different words refer to the same thing (or vice versa), but a concept in a representation may not have a corresponding concept in the other, or a single concept on one side may correspond to a complex conceptual structure on the other side.

Interpretation schemas. Two individuals (communities, groups, . . .) may use very different conceptualizations because of their different interpretation schemas. A classic example is that of paradigms in scientific communities, discussed by Thomas Kuhn in his book on *The structure of scientific revolutions* (1979). However, there are much more concrete examples in ordinary life, when people disagree on the truth of facts that look self-evident to each side. A real example is that of a company in which the marketing and the production divisions had a completely different view on where the company should invest in innovation: for the first division, it was clear that the company needed to renew the image, as the quality of their product was already good enough; for the second, it was an absolute truism that the company needed to invest in research and development, as the quality of the product is the key success factor for the company. These opposite views were built on different interpretation of what is important, good, innovative, and so on. This form of heterogeneity is extremely deep, and it involves the “identity” of people and communities).

A common strategy for achieving semantic interoperability is to create large, shared conceptual schemas (ontologies, taxonomies, classifications, knowledge maps, corporate languages) that are then used as a common reference for agents using autonomous (and typically heterogeneous) conceptualizations. Shared conceptualizations are supposed to play a role similar to the role that the reference to a common world of objects and facts plays in human communication. However, this approach seem to have into two very general problems:

Copyright © 2002, American Association for Artificial Intelligence (www.aaai.org). All rights reserved.

- on the one hand, if even the reference to a common world does not always guarantee the success of communication (i.e., communication free of ambiguities and misunderstandings), the problem can only get worse when we try to replace the common world with an abstract structure (e.g., an ontology), which – being a linguistic structure – is by definition compatible with many different (possibly contrasting) interpretations;
- on the other hand, it seems to face the resistance of many communities, that refuse to accept a third-party conceptualization of what they know. This is not surprising, if we consider the relation between knowledge and identity which we mentioned when talking about interpretation schemas. G. C. Bowker and S. L. Star, in a book on *Sorting things out: classification and its consequences* (MIT Press, 1999), observe that “[a]ny information system that neglects use and user semantics is bound to trouble down the line – it will become either oppressive or irrelevant”.

The aim of this workshop is to investigate an alternative approach to semantic interoperability, based on what we call *meaning negotiation* (MN). The idea of MN is that any viable approach to semantic interoperability between semantically autonomous entities (that is, entities that cannot assess semantic problems by “looking into each other’s head”, like humans or software agents) should involve (among other things) a social process of negotiating an agreement on the meaning (semantics) and on the speaker’s intention (pragmatics) of a communication. As a simple example, imagine that you ask a travel agent in Miami information about hotels in Venice. In the first place, the agent will try to establish which Venice you are talking about (e.g., Venice in Florida or Venice in Italy?), namely the semantics, and then the purpose of your trip (e.g., leisure or business), the dates (e.g., summer or winter), the price range (e.g., cheap or expensive), namely the intention. After this phase (in this case, based on questions and answers), the travel agent will try to provide meaningful suggestions. Other forms of MN may involve: existence (do black holes exist?); appropriateness of a classification (is the attack to the Twin Towers to be classified as a “terrorist attack” or as a “declaration of war?”); relationships between concepts (is “artificial intelligence” a sub topic of “computer science” or a branch of “cognitive psychology?”); ambiguity (under what conditions “bug” and “error” can be treated as synonyms?); and so on.

We’d like to observe that any approach based on MN requires to address to two different problems:

1. *semantic coordination*, namely the problem of reaching an agreement on the meaning of an expression when (i) agents have a common interest in achieving such an agreement, but (ii) there are many possible solutions to the problem, and thus the selection of one of these solutions can be problematic;
2. *semantic negotiation*, namely the problem of reaching an agreement on the meaning of an expression when (i) an agreement is valuable for all agents, but (ii) agents have conflicting preferences over which solution should be selected, so that every agreement implies that at least someone has to concede to some extent to other agent.

Consider a situation in which two agents use different names (e.g., “car” and “automobile”) for the same concept. If they want to exchange meaningful information, they need to find a mapping between their representations. In principle, this problem has many possible solutions, as each agent could imagine many possible mappings. The goal of any strategy of semantic coordination is precisely to narrow down the number of possible mappings, ideally to one. If they can reach such an agreement (e.g., the first agent records somewhere that there is a semantic mapping between “car” on its side and “automobile” on the other side, and vice versa for the second agent), then communication will be successful, as their use of the two words is sufficiently coordinated.

However, in other situations, such a coordination agreement is not enough. For example, we can imagine a situation in which two IT managers need to decide which between two schemas is more appropriate for an employee database after their two companies merged. Of course, the two managers will have conflicting preferences, as changing a database schema can require a massive re-implementation of a large number of applications. In this case, the disagreement cannot be explained simply in terms of coordination, as there are different stakes involved (e.g., assessing what is a legitimate reaction to an event, spending time and money to rebuild a database and possibly many related applications), and thus the agents involved have different preferences. This situation requires some strategy of semantic negotiation. The relative strength of agents, concerns of fairness, and the possibility of voluntary disagreement have to be taken into account to get an appropriate picture of the nature of negotiation problems.

As game theory made clear, coordination and negotiation are different types of interaction (although any negotiation problem has a peculiar coordination component in it). Different concepts of solution apply to them, as well as different strategic concerns emerge from the observation of human behavior.

This workshop is an attempt to bring together researchers and practitioners who have a research interest on MN from an interdisciplinary perspective. The problem is attacked using different conceptual and technological tools, and with different motivations in mind. Contributions cover a wide range of interaction patterns, from strict bargaining to situations where coordination features seem to prevail. The contributions come from several different research areas, including multi-agent systems, knowledge representation, databases, natural language processing, machine learning, game theory, philosophy of language, cognitive science, organization and management sciences. They cover many important application areas, such as the semantic web, knowledge management, human-computer interaction, computer supported cooperative work (CSCW). However, we hope this will be only the beginning of an enterprise that, in our view, can have very exciting developments, both from a theoretical and a practical point of view.