

10. Conclusion

Even in a book length account of ontological semantics, many very important issues had to be left out. Most glaringly, we have not described in any detail the processes of text generation and of various types of reasoning in support of human-computer dialog applied to a variety of information processing applications. Versions of the text generator have been developed in the Dionysus (Nirenburg *et al.* 1989) and Mikrokosmos (Beale and Nirenburg 1995, 1996; Beale *et al.* 1997, 1998) implementations of ontological semantics. Reasoning capabilities were developed, for instance, in the Savona system (Nirenburg 1998a).

In text generation, ontological semantic apparatus helps to choose what to say at the first stage of generation; offers a rich search space for lexical realization options; and provides a convenient control structure (Beale 1997) for combining realization constraints from semantic and non-semantic sources. Ontological semantics can support generation not only of text but also of text augmented by tables, diagrams, pictures and other illustrative material.

The richness of the knowledge content in the ontology and the fact database opens possibilities for enhancing automatic reasoning. The inclusion of complex events in the ontology promises better results in story recognition, in planning dialog responses and in determining the goals of the text producer. The general idea is not new: it ascends to the script processing efforts of the 1970s and 1980s (Schank 1975, 1981). However, ontological semantics not only describes complex events at a realistic, non-toy grain size but also includes sufficiently fine-grained knowledge content for all the other components needed for processing.

However large the content of all the static and dynamic knowledge sources in ontological semantics may seem, it is quite appropriate to multiply it, with changes, and encase each such complete semantic apparatus in a model of an intelligent agent. It then becomes possible to carry out application-driven experimentation of communication within a society of such artificial intelligent agents or between such agents and people. Clearly, any, even limited, success in this area promises substantial practical benefits in various applications, notably including Internet-related activities. Once again, the idea as such is not new: people have already talked about avatars and other personal agents (for the latest, see, for instance Agents-98; Agents-99; Cassell and Vilhjálmsón 1999) to help people with their information needs. It is, once again, the detail of content in ontological semantics that promises to bring such applications to a new level of quality.

To support reasoning by intelligent agents, their knowledge bases must include not only their own ontology and Fact DB (we downplay here the role of the language-oriented resources in this process) but also their impressions of the ontologies and Fact DBs of their interlocutors. This capability is important in recognition of goals and plans of other agents. Thus, if Agent A relates to Agent B that it has finished editing, Agent B, using its ontology for Agent A, recognizes the complex event of which this instance of editing is a part and then can project what the next activity of Agent A should be. Usually in applications this situation is simplified by assuming the same ontology for each agent involved. The ontologies of agents A and B may, in fact, differ. A particular complex event can be simply absent in one of the Agents' ontologies. Also, one agent's notion of the ontology and the Fact DB of the other agent may be inaccurate. Many more discrepancies are possible, all resulting in wrong conclusions.

A natural, though possibly spurious issue here is that of infinite regress of models of the ontologies of others' ontologies of oneself, etc. While some such knowledge may indeed be important in some applications, the complexity inherent in multiply nested ontologies inside each intelligent agent makes reasoning over them prohibitively expensive. It was most probably the realization of this fact that led Ballim and Wilks (1991) to curtail the levels of inclusion of beliefs of others in their model to no more than three turns.

Modeling intelligent agents using the ontological semantic apparatus will also facilitate general experimentation in dialog processing. Among many potential uses, ontological semantics can provide the knowledge and processing substrate for large-scale implementations of the ideas about treating the dialog situation; for example, for studying the levels of similarity between the speaker's and hearer's ontologies and Fact DBs necessary to attain acceptable levels of understanding in a dialog.

The development of ontological semantics continues in the direction of better coverage and finer grain size. The further refinement and automation of tools and the adaptable, hybrid methodological base of ontological semantics steadily bring down the concomitant acquisition costs and support improvement of the processing components. The emphasis on rich semantic content and the unique division of labor between humans and computers both in acquisition time and within applications, overcomes the common pessimism about applications based on representation and treatment of meaning with regard to the attainability of fully automatic, high quality results at a reasonable cost.