Ontological Foundations of Natural Language Communication in Multiagent Systems

Luc Schneider† and Jim Cunningham‡

†Institute for Formal Ontology and Medical Information Science, University of Leipzig
‡Department of Computing, Imperial College London

Abstract. The paper outlines a semantic ontology as a minimal set of top-level conceptual distinctions underlying natural language communication. A semantic ontology can serve as the basis for the specification of the meaning, as the logical form, of agent messages couched in natural language. It represents a general and reusable module in the architecture of multi-agent systems involving human as well as software agents. As a practical example, we will sketch a basic multi-agent system relying on natural language communication.

1 Ontology as a Basis for Multiagent Semantics

Successful communication in a multiagent system requires not only that the communicating agents share a common language, but also that they are committed to the same intended model for the semantics of this language. The *semantics* of a communication language is the theory that specifies the truth conditions of the messages embedded in the agents' speech acts.

Under the closed world assumption, a shared intended model may be specified as a subset of the Herbrand base, that is, the set of ground goals of the communication language. In this case an ontology can be regarded as the logic program whose declarative meaning (roughly, the set of ground goals deducible from it) is an intended model shared by a community of communicating agents. This is just a paraphrase of the classical definition of an ontology as the formal statement of a model specifying the shared understanding of cooperating agents (Gruber 1991, 1995).

A *semantic* ontology is a conceptualisation, common to a community of agents that understand natural language, of the categories and relations that pervade the agents' environment as a whole. It can be used to specify the logical form as the truth-functional meaning of agent messages embedded in natural language. Architecturally a semantic ontology is the most reusable component of multiagent systems involving a human-computer interface.

A semantic ontology has to reflect the wired-in conceptual framework human agents are equipped with. In (Schneider 2001), a minimal semantic ontology was drafted which drew its inspiration from two sources: the semantical analysis of natural language and philosophical accounts of the commonsense view of reality.

Indeed, Parsons' (1990) account of the semantics of verbs in terms of underlying events as well as the parts played in the latter by objects can be ideally complemented by Strawson's (1959) "descriptive metaphysics", an attempt to specify the basic entity-types of commonsense.

The paper is structured as follows. In Section 2, we present the basic conceptual distinctions that are required by a minimal semantic ontology. Events and processes have to be differentiated from the objects that participate in them; the participants of events are either physical objects or persons. In Section 3, we show how to classify different ways of participation according to the kinds of events and participants involved, giving an ontological reading of verb complementation. Section 4 sketches the role of a semantic ontology as a basis of natural language communication between agents by using a simple multi-agent architecture involving human and software agents as an example.

2 Basic distinctions in semantic ontology

Dependent and Independent Entities It seems to be a fundamental feature of the human conceptual scheme that some kinds of entities, like physical objects or persons, are considered as basic, while other types of individuals, like qualities or boundaries, are regarded to be somehow dependent on the former. According to Strawson (1959:16-17), this dependence has to be understood in terms of identification in an agent's environment. A class of particulars A (say, colours or boundaries) is identification-dependent on a class of particulars B (say, physical objects) if and only if, in order to be able to identify an instance of A, an agent has to single out an instance of B first.

The commonsense distinction between dependent and independent entities is also acknowledged by recent computational upper-level ontologies, like BFO (Smith 2002) or DOLCE (Gangemi et al. 2002). In particular, the dichotomy between objects and the characteristics dependent on them is fundamental for a semantic ontology underlying natural language communication, as it motivates the grammatical difference between nouns and adjectives. The common role of nouns is to refer to objects or kinds of objects, while adjectives usually denote attributes. Of course, there are exceptions to that rule, but nominalisations of adjectives, such as "Green" or "wisdom" seem to be recognised by speakers as exceptions to a more basic semantic rule.

Persons and Bodies Another distinction that is crucial for a semantic ontology is that between between *mental* or *private* characteristics (e.g. beliefs, intentions, desires) on the one hand, and *physical* or *public* characteristics (e.g. weight, colour) on the other hand. According to Strawson, our conceptual equipment is such as to posit the distinction between two types of spatio-temporal objects, namely *bodies*, to which only physical attributes can be ascribed, and *persons* which both mental and physical characteristics can be attributed to (Strawson 1959:102-103).

Many natural languages reflect this distinction explictly, by gender or other systems of noun classifications. We will see that the *Person/Body* dichotomy even underlies the semantical subcategorisation or complementation of verbs. Thus cognitively oriented ontologies like DOLCE (Gangemi et al. 2002) have to include the difference between agentive and non-agentive objects in their taxonomies.

Objects and Events Following Davidson (Davidson 1980), Parsons defends the view that the semantics of verbs and verb phrases implies the existence of events and processes (1990:4, 186-187): verbs may be considered to represent kinds of *processes* or *events*. However, the idea that the grammatical distinction between nouns and verbs is grounded on the ontological dichotomy of objects versus events or processes has been always intuited by natural language syntacticians (Tesnière 1959).

Objects persist through time in virtue of core characteristics that are fully present throughout their life. Processes exist in time by having different phases at different instants, except events as instantaneous boundaries of processes (Simons 1987). Strawson argues that events or processes are dependent on objects with regard to their identification (1959:39, 45-46). Objects enjoy an ontological priority over events or states. The dependence of an occurrence on an object is called *participation* in DOLCE (Gangemi et al. 2002).

3 Object Participation in Language

The different ways objects participate in occurrences (processes or events) have been studied by linguists interested in the phenomenon of verb complementation or *thematic roles*. These are partly syntactic, partly semantic relations between noun phrases and the main verb of a sentence. Thematic roles correspond to the different parts that referents play in the occurrence expressed by the verb (Parsons 1990:72-73). Table 1 shows Parsons' (1990:73-78) list of thematic roles together with their definitions and examples.

Obviously, Parsons' empirically assembled list lacks an ontological systematisation. We count three subject-related roles: *Agent*, *Experiencer* and *Performer*, where the *Person/Thing* and *Private/Public* distinctions are muddled together. In (Schneider 2001), a coherent ontological account of thematic roles and ways of participation is given according to the following lines. Firstly, we will consider as basic only those thematic roles which express mere specifications of the participation relation that are neutral as to the types of occurrences or objects involved. The result is shown in Table 2.

Secondly, *Agent*, *Experiencer* and *Performer* are defined using our four elementary thematic roles and the basic particular-types. We are here in the presence of two orthogonal oppositions:

1. *Agent* or *Experiencer* vs. *Performer*: the difference between personal and non-personal origins of occurrences;

Thematic Roles	**Definition**	**Example sentences**
Agent	Person initiating the event	*John* writes a book. The book is signed *by John*.
Theme	Entity affected by the event	Mary reads *a book*. *Mary* blushed at his sight.
Goal	Addressee	John gives *Mary* a rose. Anna writes a letter *to Mary*.
Benefactive	Entity to whose benefit the event occurs	Mary gave *Anne* a party. John signs a book *for Mary*.
Experiencer	Person the event is an experience of	*Mary* sees a rose. *John* thinks about Mary.
Instrument	Thing the event is accomplished with	John opens the letter *with a knife*.
Performer	Thing initiating the event	*The knife* opened the letter.

Table 1. Parson's Classification of Thematic Roles

2. *Agent* or *Performer* vs. *Experiencer*: the difference between a physical and a mental occurrence of which the object is an origin.

To clean this orthogonal classifications up, we first define a new thematic role, namely *Initiator*, and redefine *Performer*, as restrictions of the *Origin* role to the object-types *Person* and *Body* respectively. An object x is an *initiator* of an occurrence y if and only if x is an origin of y and x is a person. An object x is a *performer* of an occurrence y if and only if x is an origin of y and x is a body.

The thematic roles of *Agent* and *Experiencer* are then characterised as specifications of the *Initiator*-roles. Indeed, if x is an initiator of an occurrence y, then x is an *agent* of y if and only if y is a public or physical occurrence; x is an *experiencer* of y if and only if y is a mental or private occurrence of x.

Thematic Roles	**Definition**	**Example sentences**
Origin	Entity initating the event	*John* writes a book. *A stone* hits the window. The book is signed *by John*. The window was hit *by a stone*.
Theme	Entity affected by the event	Mary reads *a book*. *Mary* blushed at his sight.
Addressee	Entity the event is directed to	John gives a rose *to Mary*. Mary gives water *to her flowers*.
Benefactive	Entity to whose benefit the event occurs	Mary gave *Anne* a party. John signs a book *for Mary*.

Table 2. Revised Classification of Basic Thematic Roles

Thus by using basic ontological distinctions, we can transform a flat unsystematized list of thematic roles into a reasoned taxonomy.

4 A Proof of Concept

The role of a semantic ontology with respect to natural language understanding in multi-agent systems has been exemplified in (Schneider 2001) by implementating, as a simple proof of concept, a reasoning agent as a server capable of processing natural language queries from multiple human operated clients. Concretely, this reasoning agent is able to engage in a game of challenges and answers: opponents send natural language assertions to be parsed, proved or disproved, the agent justifying its answers by indicating the respective logical form (meaning) or proof established on the basis of a semantics/ontology shared with the human opponent.

A way to realise such a system is to implement it as a client-server architecture, the server being the reasoning agent and the client(s) operated by human users. The server spawns off a new thread for each client, thus allowing peer-to-peer communication. Multi-agency is thus not merely implemented by pairing off a single program with a single human, but actually involves multi-threading and inter-thread communication. This architecture has been implemented in Qu-Prolog, a distributed and concurrent version of Prolog (Clark, Robinson and Hagen 1999; Robinson 2000).

By spawing a new thread or agent at each client's request, the reasoner server becomes the central node in a multi-agent system of communicating human and non-human peers. The content of the humans' messages are declarative natural language sentences whose meaning, i.e. logical form, mirrors the everyday conceptual framework of intelligent primates. Shared understanding is made possible by arranging that the software agents have the same ontology, i.e. set of fundamental conceptual distinctions, as their human partners.

This ontology is the basis for the semantics of the natural language fragment used by the humans to communicate with their non-human peers. Sharing this ontology as a part of their knowledge base, the software agents have the capability of parsing and proving , i.e. of understanding and reasoning upon the assertions submitted to them. The logical form and proof computed by a thread of the reasoning server reflects the semantic and ontological intuitions of the human operators.

The primitive and defined predicates of the semantic ontology, i.e. the basic particular-types $Person$ and *Body* as well as the thematic roles discussed in the previous section, are used directly in parsing the natural language sentences submitted by the human users. The parser is basically a logic grammar translating into action the semantical analysis of verbs and sentences in terms of underlying events and thematic roles. The thematic roles occurring in the entries of verb meanings in the parser's lexicon are declared in the semantic ontology, which constitutes a separate module of the agent's knowledge base.

As an illustration, we describe a sample run of the system presented in Schneider 2001 (Figure 1). After connecting to the server, the user requests the proof of the natural language sentence: "joan liked marcel" by the server thread spawned for that purpose. First, the server thread parses this sentence into its logical form expressing that there exists an occurrence, which is a liking, whose *exp*eriencer is Joan and whose *th*eme is Marcel.

Second, the server thread proves this first-order logical formula using its knowledge base, displaying the steps of the deduction. Indented steps indicate backtracking triggered successively on the conjuncts in the body of the definitions of *Experiencer*, *Performer* and *Private* respectively. These definitions are not part of the particular domain-related knowledge of the agent, but belong to the semantic ontology as a separate module.

```
| ?- reasoner_prove.
> joan liked marcel

exists : (_34A , ((liking(_34A) , (exp(_34A, joan) , th(_34A, marcel)))))

     liking(o(20))
     exp(o(20), joan)
          perf(o(20), joan)
               or(o(20), joan)
               person(joan)
          private(o(20))
               liking(o(20))
     th(o(20), marcel)

yes
```

Fig. 1. A sample run

By constructing the semantic ontology as a separate module in the architecture of a reasoning agent, two essential goals for agent implementations are achieved: scalability, reusability and maintainability. By storing the declarations and definitions of the various ontological types and roles outside of the lexicon, the latter can be scaled down in size, thus enhancing the efficiency of the parser. As a component of its own, a semantic ontology is easier to share between applications and to reuse in various contexts involving different parsing technologies. Finally, a semantic ontology as a distinct module is trivially easier to maintain without the need of modifying other components of the multi-agent architecture.

5 Conclusions

The aim of this paper has been to outline a minimal semantic ontology, a set of high-level concepts that can serve as a basis for specifying the logical form of agent messages using natural language. Its main inspiration comes from the semantical analysis of natural language, as well as philosophical accounts of the commonsense view of the world. A semantic ontology may be put to two uses: to define the fundamental concepts necessary for agents to communcate and to reason, as well as to contribute to the computational analysis of natural language. The two uses can be combined in a multi-agent system involving humans and thus recurring to natural language communication. A simple instance of a multi-agent architecture based on natural language communication that illustrates these both uses has been sketched at the end of this paper. We emphasised the advantages in terms of scalability, reusability and maintainability of having a semantic ontology as a separate module in a multi-agent architecture.

Acknowledgements This paper is partly based on work supported by the Alexander von Humboldt Foundation under the auspices of its Wilhelm Paul Programme.

References

Clark, K., Robinson, P. and Hagen, R. 1999. Multi-Threading and Message Communication in Qu-Prolog. Technical Report 99-41, Software Verification Research Center, University of Queensland.

Davidson, D. 1980. "The Logical Form of Action Sentences". In Davidson, D., *Essays on Actions and Events.* Oxford: Oxford University Press, 105-148

Grenon, P. 2003. "The Formal Ontology of Spatio-Temporal Reality and its Formalization". *AAAI Spring Symposium on the Foundations and Applications of Spatio-Temporal Reasoning 2003.*

Gruber, T. 1991. "The Role of Common Ontology in Achieving Sharable, Reusable Knowledge Bases". *Principles of Knowledge Representation and Reasoning: Proceedings of the Second International Conference.* Morgan Kaufmann, San Mateo/CA.

Gruber, T. 1995. "Toward Principles for the Design of Ontologies Used for Knowledge Sharing". *International Journal of Human and Coputer Studies. Special Issue: Formal Ontology, Conceptual Analysis and Knwoledge Representation.*

Gangemi A., Guarino N., Masolo C., Oltramari, A., Schneider L., 2002. "Sweetening Ontologies with DOLCE". *Proceedings of EKAW 2002.* Siguenza, Spain.

Parsons, T. 1990. *Events in the Semantics of English: A Study in Subatomic Semantics.* Cambridge/MA: MIT Press.

Robinson, P. 2000. Qu-Prolog 6.0 User Guide. Technical Report 00-20, Software Verification Research Center, University of Queensland.

Schneider, L. 2001. "Naïve Metaphysics. Merging Strawson's Theory of Individuals with Parson's Theory of Thematic Roles as a Basis for Multiagent Semantics". MSc. Computing Science thesis, available as *Departmental Technical Report 2002/5.* London: Department of Computing, Imperial College London.

Simons, P. 1987. *Parts. A Study in Ontology..* Oxford: Clarendon.

Smith, B. 2002. "Basic Formal Ontology". http://ontology.buffalo.edu/bfo/

Tesnière, L. 1959. *Éléments de syntaxe structurale.* Klincksieck, Paris.

Strawson, P. 1959. *Individuals. An Essay in Descriptive Methaphysics.* London: Routledge.