An Ontology-Based Utterance Interpretation in the Context of Intelligent Assistance

Emerson Cabrera Paraiso[1], Jean-Paul A. Barthès[2]

[1]Programa de Pós-Graduação em Informática – Pontifícia Universidade Católica do Paraná (PUCPR) - CEP: 80215-901 – Curitiba, PR – Brazil

[2]Université de Technologie de Compiègne (UTC) – Laboratoire Heudiasyc – Centre de Recherches - B.P. 20.529 - Compiègne - France

`paraiso@ppgia.pucpr.br, barthes@utc.fr`

Abstract. *In this paper we present an ontology-based utterance interpretation in the context of intelligent assistance. Ontologies are used for syntactic and semantic interpretation and for task representation. This mechanism is embedded in a conversational interface applied to personal assistant agents. The main goal of this approach is to offer a system capable of performing tasks through an intuitive interface, allowing experienced and less experienced users to interact with it in an easy and comfortable way.*

1. Introduction

A Personal Assistant (PA) is a specialized artificial agent that helps human users to do their daily work. The main goal of such an agent is to reduce the user's cognitive overload. In such a context, the agent displays autonomy while working with the user cooperatively as shown by Maes [Maes 1994]. In an agent-based system, as in any other software application, an appropriate user interface is crucial. Traditionally, developers propose graphics-oriented interfaces involving the use of menus, sub-menus, dialogue-boxes, and so on. Often this approach is inappropriate or at least not very appealing, leading to an assistance of poor quality [Paraiso and Barthès 2006]. An attractive solution is conversational interfaces. Conversational interfaces as defined by Kölzer [Kölzer 1999], let users state what they want in their own terms, just as they would do, speaking to another person. One of the most difficult tasks in implementing a conversational interface is to interpret utterances and understand their meaning. To do that, we are using ontologies. The ontologies play a key role at the semantic interpretation time since the meaning of utterances can be inferred by looking for concepts and their attributes. The use of ontologies for representing domain knowledge and for supporting reasoning is becoming widespread. The ontologies however, may also be used for facilitating the interaction between user and PA. The concepts and their properties are organized to map the world but also to help processing natural language. In this paper, we present an ontology-based utterance interpretation in the context of intelligent assistance.

This paper begins by describing how ontologies are used for syntactic and semantic interpretation and for task representation. We present some examples for better understanding. Finally, we offer a conclusion and indicate some perspectives.

2. Interpreting Users Utterances

Utterances are statements presented by users to the PA. In dialogue, each turn by a speaker may be considered an utterance. In our approach, the process of interpreting an utterance is carried out in two steps: a) parsing and syntactic analysis; and b) semantic interpretation (ontology application). The parsing algorithm works top-down, replacing each utterance stem with its syntactic category (verb, noun, adverb, etc) with the help of lexicon files and a set of grammar rules. In order to reduce the interaction and, consequently avoid wasting time, we limited the space of dialogue utterances to directive speech act classes [Searle 1975]—inform, request, or answer—since such classes define the type of expected utterances in a master-slave relationship. A speech act is an act that a speaker performs when making an utterance. The idea is to do something by looking for the related acts in the utterance.

In our applications, a typical utterance could be: "*Look for a document on agents in the database*." According to our taxonomy this is an order utterance and can be processed by the grammar rules. If a sentence is not well formed, or if it is out of the domain, then it is classified as a nonsensical utterance.

2.1 Using the Ontologies

In the context of an open conversation, the problem of understanding is complex, demanding a well structured knowledge base. Domain knowledge is used here to further process the user's statements and for reasoning. To this effect, we are using a set of task and domain ontologies, separating domain and task models for reasoning. As suggested by Allen [Allen et al. 2001], the approach is interesting for domains where task reasoning is crucial. Ontologies are a trend in the dialogue community. The key components that make up an ontology are a vocabulary of basic terms and a precise specification of what those terms mean [Guarino 1998]. Ontologies play two main roles in our PA: a) they help interpreting the context of messages sent by others agents or by the user (utterances); and b) they keep a computational representation of knowledge useful at inference time. The ontologies may also facilitate the process of semantic interpretation, supplying the parser with linguistics elements, like noun synonyms, or hyponyms/hyperonyms.

The approach to semantic interpretation presented here is based on the notion that the meaning of utterances can be inferred by looking for concepts and their attributes. More precisely, the module responsible for applying the ontology to the utterance searches for domain concepts and the list of verbs that indicates the task to be executed. The corresponding keywords are concepts of the ontology directly related to a list of actions. We believe that this approach is ideal for applications where the domain is well known and restricted. In this paper, we reduced the ontologies to a bare minimum for explaining the semantic interpretation mechanism. The concepts and their properties were organized to map the world but also to help processing natural language (e.g. by adding a list of applicable actions to each concept of the ontology).

2.2 Some Examples

To illustrate how the mechanism works, consider the utterance:

USER: Could you search all articles on Agents?

A very simple piece of ontology is shown Figure 1a (we used Protégé [Gennari et al. 2002] for an easier representation), describing concepts that model a project. A project, according to the ontology, may have different types of documents, an address book, an agenda, a list of members and a list of actions related to each concept. Each concept may have some attributes (Figure 1d). Note that a set of actions (e.g., read, list, erase, as shown in Figure 1b) may be applied to each concept, as shown in Figure 1c.
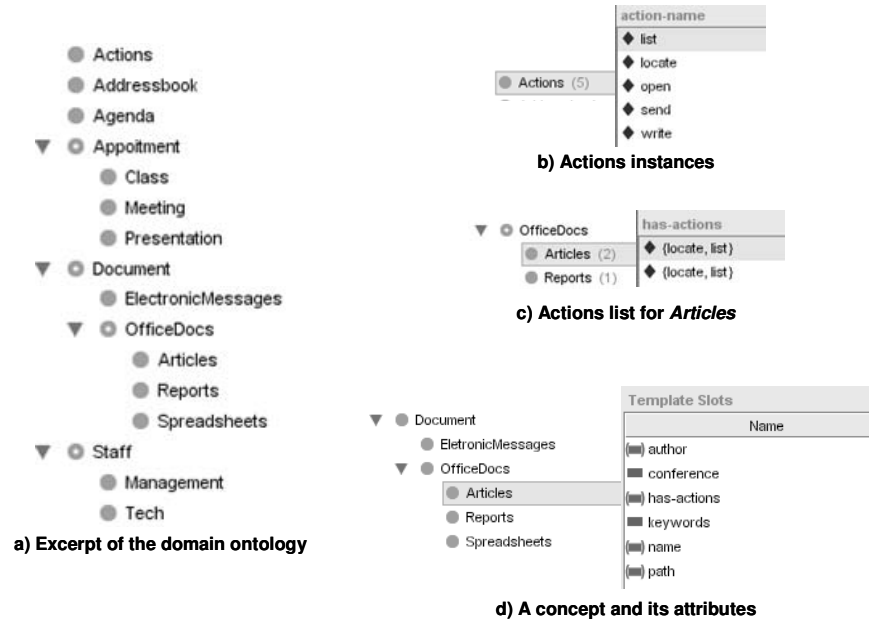

b) Actions instances

c) Actions list for *Articles*

a) Excerpt of the domain ontology

d) A concept and its attributes

Figure 1. An excerpt of the domain ontology *Project*

To interpret the given utterance, first the parser checks the context of the input, verifying that it is a question related to the domain. To do so, it uses the domain ontology and the lexicon. Since it is a question and since it is related to the application domain, the Grammar Verification module returns a matrix containing the list of tokens and their syntactic classification. By looking up the tokens in the ontology, it finds that the token search is an action (Figure 1c). Note that it uses a list of synonyms (e.g. "locate", "search" or "find" are synonyms in this sense). It finds also that *articles* is a concept and *Agents* is one of its properties (subject). A formal representation is then produced and stored in memory.

The PA's knowledge base is also composed of a list of task descriptions. A task description models a specific service provided by an agent. Each task has five parameters and each parameter has several fields. A task is performed by a service agent when it receives a request from another agent. The Dialogue Manager module will push a task onto the stack of tasks when an utterance related to the task is given. After identifying the task, the PA tries to fill all task parameters asking the user to give more information if needed. The dialogue manager can handle many tasks simultaneously (even tasks of the same type). Let us take an example. The user says:

USER: I need to send an email to Mike Palmer.

After parsing and semantic analysis, the Dialogue Manager is able to start a new task, since it is related to the domain (according to our first ontology presented in Figure

1). The task *To Send an Electronic Message* has some parameters to be filled before the agent is able to execute it. One of the parameters may be the subject of the message. Since the given utterance does not contain this information, the Dialogue Manager will request it from the user, asking him/her the question defined in the appropriate question field. The Dialogue Manager changes the task status to pending and waits for a response from the user. When all fields are filled, the Dialogue Manager sends the task for execution.

3.Conclusions and Future Work

In this paper, we presented an ontology-based utterances interpretation approach. The effectiveness of this architecture may be achieved by applying special designed domain and task ontologies. They have been used to aid the semantic interpretation process with success. Thanks to them, impossible requests, such as those out of context, are handled easily.

Following this work, possible future applications include the improvement of the agent's memory, since ontological instances used to interpret utterances may also be used for reference treatment. Another important direction for future work consists in allowing users to access their assistant via different means, like telephones or personal digital assistants (PDA). The user's mobility may impose interesting constraints to be tackled by future versions of our speech interface.

References

Allen, J., Ferguson, G. and Stent, A. (2001) "An Architecture for More Realistic Conversational Systems", In: Proceedings of Intelligent User Interfaces 2001 (IUI-01). Santa Fe - NM.

Gennari, J., Musen, M. A. , Fergerson, R. W., Grosso, W. E., Crubézy, M., Eiksson, H. Noy, N. and Tu, W. (2002) "The Evolution of Protégé: An Environment for Knowledge-Based Systems Development". Report available at: http://smi.stanford.edu/pubs/SMI_Abstracts/SMI-2002-0943.html.

Guarino, N. (1998) "Formal Ontology in Information Systems", In: Proceedings of FOIS'98, Italy, IOS Press, p. 3-15.

Kölzer, A. (1999) "Universal Dialogue Specification for Conversational Systems", In: Proceedings of IJCAI 99 – Workshop on Knowledge and Reasoning in Practical Dialogue Systems.

Maes, P. (1994) "Agents that Reduce Work and Information Overload", In: Communications of ACM, Vol. 37, No.7, ACM Press, pp. 30-40.

Paraiso, E. C. and Barthès, J.-P. A. (2006) "An Intelligent Speech Interface for Personal Assistants in R&D Projects", Expert Systems with Applications, Volume 31, Issue 4, p. 673-683.

Searle, J. R. (1975) "A Taxonomy of Illocutionary Acts", In: Proceedings of Language, Mind and Knowledge, Vol. 7, University of Minnesota Press, p. 344-369.