Dynamic Use of Ontologies in Dialogue Systems

Joana Paulo Pardal

Department of Information Systems and Computer Engineering
Instituto Superior Técnico, Technical University of Lisbon
Lisbon, Portugal
`joana@l2f.inesc-id.pt`

Abstract

Most dialogue systems are built with a single task in mind. This makes the extension of an existing system one of the major problems in the field as large parts of the system have to be modified. Some recent work has shown that ontologies have a role on the domain knowledge representation as the knowledge collected in an ontology can be used in all the modules. This work aims to follow the footsteps of the use of ontologies in dialogue systems and take it further as the current state of the art only uses taxonomical knowledge.

1 Introduction

At the present time, the Spoken Language Systems Lab (L^2F) integrates a project in the "House of the Future" at the Portuguese Communications Foundation. The house has a spoken dialogue system (Mourão et al., 2004) based on TRIPS architecture (Allen et al., 2005) where a virtual butler named "Ambrósio" helps the user in daily tasks that deal with devices and services, through speech commands. Whenever clarification is needed, further dialogue is entailed. To act in response to the user, the system needs to know which devices are connected, which services are available and what actions can be performed. Currently, this information is stored for each service or device: the available operations, the needed parameters and the possible values for each one. This kind of architecture is very common in the

field. Nevertheless it's still hard to extend an existing system because it's always necessary to adapt lots of features in the system.

Recent work from Filipe (2006) has enhanced the access to the services and abstracted the database view in order to create an Application Programming Interface (API). The main contribution of that work is a Domain Knowledge Manager (DKM) advisor service, which suggests the best task-device pairs to satisfy a request. Additionally, a DKM recognizer service to identify the domain concepts from a natural language request is proposed. A hybrid approach is used to design ubiquitous domain models to allow the dialogue system to recognize the available devices and tasks they provide on-the-fly.

But more work is still needed to ease the dynamic configuration of dialogue systems and to deal with a set of arbitrary plug-and-play devices. The main goal of this work is to pursue the work done by Filipe.

2 State of the art

This work encompasses knowledge and techniques from two different areas: dialogue systems and ontologies. This work has to deal with the challenges from all these areas.

2.1 Dialogue Systems

Since the 1980s, the Natural Language Processing community has used spoken dialogue systems as a case study (Colea et al., 1997). This option is explained by the simplicity that comes from the treatment of restricted domains. The multidisciplinarity involved is one of the richnesses of this field as

it brings together people from several communities like signal processing – for speech recognition (Jurafsky and Martin, 2000) and synthesis (Huang et al., 2001); artificial intelligence – for interpretation of the spoken utterances (Allen, 1987); and software engineering – for more efficient architectures (McTear, 2002). But the complexity of these systems makes them expensive to develop (Allen et al., 2000) and difficult to adapt to new types of users, services, languages and scenarios (Turunen and Hakulinen, 2003).

With the proliferation of databases, some work has been done to take advantage of the knowledge structure and organization to dynamically extend existing systems to new domains, devices and services.

2.2 Ontologies

Ontologies aim at capturing static domain knowledge in a generic way and providing a commonly agreed understanding of a given domain. The main purpose is to share and reuse that knowledge across applications. The field of Ontologies appeared in the 1990s (Gruber, 1993), but only lately has been perceived as more valuable, as some effective results are being achieved with their use, reuse and sharing. Being so, an ontology is a formalized shared specification of a conceptualization. Mainly, a domain ontology collects the relevant concepts of a domain and the relations between them. An ontology usually also represents some formal restrictions verified in the domain. Therefore, ontologies usually have three types of entities: classes, relations, and axioms.

Currently the main challenges in this area include the definition of a clear building process (Pinto and Martins, 2004), automatic learning of ontologies (Maedche and Staab, 2004), transparent access to information (Gil et al., 2005) and efficient inference based on the available knowledge (Baader et al., 2003). Some work has been done where databases and other legacy knowledge sources are replaced by ontologies in different types of domains with success (Grau et al., 2005).

2.3 Use of Ontologies in Dialogue Systems

Separating the domain knowledge from the language features of the spoken dialogue systems has pro-

ven to reduce the complexity of a dialogue system's components. Moreover, if the domain knowledge is already available, reusing it is crucial to reduce the effort needed to build a new dialogue system or to extend an existing one into a new subject. Some recent work has shown the advantages of the use of Ontologies for these tasks.

Milward and Beveridge (2003) maintain that the ontology-based dialogue system for home information and control provides a dynamically reconfigurable system were new devices can be added and users can subscribe to new ones; asynchronous device input is allowed; unnatural scripted dialogues are avoided; and a flexible multimodal interaction for all users including the elderly and the disabled is provided. Also, the recognition, interpretation, generation and dialogue management are more flexible as the knowledge coded on the ontology can be used dynamically.

Flycht-Eriksson (2004) argues that the separation of the dialogue management from the domain knowledge management is crucial to reduce the complexity of the systems and enhance further extensions.

Both these works focus on the IS-A and PART-OF relations to solve under/over specification. This is helpful in medical-related dialogue systems that need taxonomical knowledge of the domain. Using more relations is still a challenge as the complexity increases.

3 Main goals

The main goal of this project is to enhance spoken dialogue systems to make them more general and domain-independent. This means that knowledge should be introduced in the system more easily and transparently. To do this, the dialog management should be separated from the domain knowledge management. This should be done not only by assigning a system module to it (the service manager) that has to be adapted to each domain, but, additionaly, by defining the kind of domain knowledge needed and creating an abstraction to represent it. For example, the dialogue system needs to know the possible words in the next expected response from the user and that depends mainly on the domain. This separation eases the creation of mechanisms to treat the common dialogue phenomena. A library

for these phenomena should be reused in dialogue systems across all domains.

Contributions from the ontologies field will be explored in regard to knowledge manipulation in a generic spoken dialogue system. As said before, some work has been done in the field but, at least for now, most of the work is reduced to the hierarchical knowledge (classes and IS-A relations) and under/over specification (PART-OF relations) that usually are represented on the ontologies. The extra-taxonomical knowledge is still being ignored but should be considered as that is the main richness of ontologies.

The most interesting topic is whether ontologies can enrich a spoken dialogue system and be used by it in such a way that the system can abstract the knowledge source thus allowing the system to focus only on dialogue phenomena and rather than the architecture adaptation that has to be done in order to include new domains.

The definition of the dialogue system as the instantiation of a spoken dialogue system will be explored after the existing dialogue systems and ontologies have been studied and categorized according to the tasks they perform and the used knowledge sources.

4 Completed Work

An ontology on the cooking domain has been built (Ribeiro et al., 2006; Batista et al., 2006). This ontology still hasn't been used but it will be included in our dialogue systems to provide help during the execution of a recipe. Currently an undergraduate student is enriching this ontology with a collection of recipes automatically extracted from text.

Also, a first prototype version of a cooking butler has been implemented. It lets the user choose from a list of recipes one to be dictated to him. Forward and rewind commands are available. This work is still preliminary as it doesn't use any ontology. It was done by two undergraduate students as a proof of concept that our current system can be extended to a dictating task.

5 Future directions

Since the PhD is still on going, lots of work is yet to be done. The next step to achieve the main goal of this work is to study the existing dialogue systems with emphasis on the performed tasks and the used knowledge sources. Beyond the simple enumeration of all the published systems, the aim is to create a categorization of dialogue systems according to the tasks they allow and to the type of knowledge they use independent of the used knowledge representation primitives (classes, relations and axioms).

5.1 Tasks to be performed

- A survey on the existing ontologies according to the coded information: classes, relations and axioms.

- Exploratory work on how to manage the domain knowledge transparently, focusing on the integration of ontologies in dialogue systems.

- Arrange the current architecture to consider not only the TRIPS architectural proposal, but the contributions coming from the ontological field. The separation of the dialogue manager in two modules should be considered here: one module for the dialogue features independent from the domain and other for the domain knowledge management.

- Adapt the existing L^2F's spoken dialogue system to the identified requirements in order to use domain knowledge from an ontology.

- Use the proposed methodology to include a cooking ontology on the L^2F's dialogue system to extend it to new domains.

- Include ontologies from different domains. An entertainment (Theatre, Movies, etc) domain ontology is being build.

5.2 Intelectual Contributions

- Classification of the existing dialogue systems according to the type of information they need and use;

- Classification of the used ontologies in dialogue systems according to the information coded and the used classes, relations and axioms;

- Propose an architecture where the contribution of each module is clearer and where the information flows both forward and backward;

- Propose a methodology for the integration of ontologies into general dialogue systems according to their classification;

- Integration of a cooking ontology into the existing dialogue system;

- Integration of another ontology into another dialogue system (from UoR).

References

James Allen, Donna Byron, Myroslava Dzikovska, George Ferguson, Lucian Galescu, and Amanda Stent. 2000. An architecture for a generic dialogue shell. *Natural Language Engineering*, 6(3).

James Allen, George Ferguson, Mary Swift, Amanda Stent, Scott Stoness, Lucian Galescu, Nathanael Chambers, Ellen Campana, and Gregory Aist. 2005. Two diverse systems built using generic components for spoken dialogue (recent progress on TRIPS). In Ann Arbor, editor, *Proc. of the Interactive Poster and Demonstration Sessions at the 43rd Annual Meeting of ACL*, pages 85–88, Michigan, USA.

James F. Allen. 1987. *Natural Language Understanding*. Benjamin Cummings, 2nd edition.

Franz Baader, Diego Calvanese, Deborah McGuinness, Daniele Nardi, and Peter Patel-Schneider, editors. 2003. *The Description Logic Handbook: Theory, Implementation, and Applications*. Cambridge University Press.

Fernando Batista, Joana Paulo Pardal, Paula Vaz Nuno Mamede, and Ricardo Ribeiro. 2006. Ontology construction: cooking domain. Technical report, INESC-ID, Lisboa, Portugal.

Ron Colea, Joseph Mariani, Hans Uszkoreit, Giovanni Batista Varile, Annie Zaenen, Antonio Zampolli, and Victor Zue (editors), editors. 1997. *Survey of the State of the Art in Human Language Technology*. CSLU, CMU, Pittsburgh, PA.

Porfírio Pena Filipe and Nuno J. Mamede. 2006. A domain knowledge advisor for dialogue systems. In *International Joint Conference IBERAMIA/SBIA/SBRN 2006 – 4th Workshop in Information and Human Language Technology*.

Annika Flycht-Eriksson. 2004. *Design and Use of Ontologies in Information-providing Dialogue Systems*. Ph.D. thesis, School of Engineering at Linköping University.

Yolanda Gil, Enrico Motta, Richard Benjamins, and Mark Musen, editors. 2005. *The Semantic Web – 4th ISWC*, volume 3729 of *LNCS*. Springer, Ireland.

Bernardo Cuenca Grau, Ian Horrocks, Bijan Parsia, and Peter Patel-Schneider, editors. 2005. *What Have Ontologies Ever Done For Us: Potential Applications at a National Mapping Agency*, volume 188.

Thomas R. Gruber. 1993. A translation approach to portable ontology specifications. *Knowledge Acquisition*, 5(2):199–220.

Xuedong Huang, Alex Acero, and Hsiao-Wuen Hon. 2001. *Spoken Language Processing: A Guide to Theory, Algorithm, and System Development*. Prentice Hall.

Daniel Jurafsky and James H. Martin. 2000. *Speech and Language Processing: An Introduction to Natural Language Processing, Speech Recognition, and Computational Linguistics*. Prentice-Hall.

Alexander Maedche and Steffen Staab, 2004. *Handbook on Ontologies*, chapter Ontology learning. International Handbooks on Information Systems. Springer.

Michael McTear. 2002. Spoken dialogue technology: enabling the conversational interface. *ACM Computing Surveys*, 34(1):90–169.

David Milward and Martin Beveridge. 2003. Ontology-based dialogue systems. In *3rd Workshop on Knowledge and Reasoning in Practical Dialogue Systems – 18th IJCAI03*.

Márcio Mourão, Renato Cassaca, and Nuno Mamede. 2004. An independent domain dialog system through a service manager. In *Proc. of 4th Intl. Conf. EsTAL*, pages 161–171. Springer-Verlag.

H. Sofia Pinto and João Pavão Martins. 2004. Ontologies: How can they be built? *Knowledge Information System*, 6(4):441–464.

Ricardo D. Ribeiro, Fernando Batista, Nuno J. Mamede Joana Paulo Pardal, and H. Sofia Pinto. 2006. Cooking an ontology. In *12th Intl. Conf. on AI: Methodology, Systems, Applications*, volume 4183, pages 213–221, Berlin.

Markku Turunen and Jaakko Hakulinen. 2003. Jaspis[2] - an architecture for supporting distributed spoken dialogues. In *Proc. of Eurospeech*, pages 1913–1916.