A Discourse and Dialogue Infrastructure for Industrial Dissemination

Daniel Sonntag, Norbert Reithinger, Gerd Herzog, and Tilman Becker

German Research Center for AI (DFKI)
Stuhlsatzenhausweg 3, 66123 Saarbruecken, Germany

Abstract. We think that modern speech dialogue systems need a prior usability analysis to identify the requirements for industrial applications. In addition, work from the area of the Semantic Web should be integrated. These requirements can then be met by multimodal semantic processing, semantic navigation, interactive semantic mediation, user adaptation/personalisation, interactive service composition, and semantic output representation which we will explain in this paper. We will also describe the discourse and dialogue infrastructure these components develop and provide two examples of disseminated industrial prototypes.

1 Introduction

Dialogue system construction is a difficult task since many individual natural language processing components have to be combined into a complex AI system. Theoretical aspects and perspectives on communicative intention have to meet the practical demands of a user machine interface—the speech based communication must be natural for humans, otherwise they will never accept dialogue as a proper means of communication with machines. Over the last several years, the market for speech technology has seen significant developments [14] and powerful commercial off-the-shelf solutions for speech recognition (ASR) or speech synthesis (TTS). Even entire voice user interface platforms (VUI) have become available. However, these discourse and dialogue infrastructures have only moderate success so far in the entertainment or industrial sector. This is the case because a dialogue system as complex AI system cannot easily be constructed. Additionally, the dialogue engineering requires many customisation works for specific applications.

We implemented a new discourse and dialogue infrastructure for semantic access to structured and unstructured information repositories for industrial applications. In this paper, we basically provide two new contributions. First, we provide architectural recommendations of how new dialogue infrastructures may look like. Second, we discuss which components are needed to convey the requirements of dialogical interaction in multiple use case scenarios and with multiple interaction devices. To meet these objectives, we implemented a distributed, ontology-based, dialogue system architecture where every major component can be run on a different host, increasing the scalability of the overall system. Thereby, the dialogue system acts as the middleware between the clients and the backend services that hide complexity from the user by presenting aggregated ontological data. We also implemented the attached dialogue components within the architecture. This paper is structured as follows. First we will discuss related work and the basic dialogue architecture. This will be followed by the discussion of individual dialogue tasks and components. Finally, we will discuss two industrial dissemination prototypes and provide a conclusion.

Related Work The dialogue engineering task is to provide dialogue-based access to the domain of interest, e.g., for answering domain-specific questions about an industrial process in order to complete a business process. Prominent examples of integration platforms include OOA [11], TRIPS [1], and Galaxy Communicator [19]; these infrastructures mainly address the interconnection of heterogeneous software components. The W3C consortium also proposes intermodule communication standards like the Voice Extensible Markup Language VoiceXML[1] or the Extensible MultiModal Annotation markup language EMMA[2], with products from industry supporting these standards[3]. In addition, many systems are available that translate natural language input into structured ontological representations (e.g., AquaLog [10]), port the language to specific domains, e.g., ORAKEL [3], or use reformulated semantic structures NLION [16]. AquaLog, e.g., presents a solution for a rapid customisation of the system for a particular ontology; with ORAKEL a system engineer can adapt the natural language understanding (NLU) component [4] in several cycles thereby customising the interface to a certain knowledge domain. The system NLION uses shallow natural language processing techniques (i.e., spell checking, stemming, and compound detection) to instantiate an ontology object. The systems which integrate sub-tasks software modules from academia or industry can be used off the shelf. However, if one looks closer at actual industrial projects' requirements, this idealistic vision begins to blur mainly because of software infrastructure or usability issues. These issues are explained in the context of our basic architecture approach for industrial dissemination.

Basic Architecture We learned some lessons which we use as guidelines in the development of basic architectures and software infrastructures for multimodal dialogue systems. In earlier projects [27, 17] we integrated different subcomponents into multimodal interaction systems. Thereby, hub-and-spoke dialogue frameworks played a major role [18]. We also learned some lessons which we use as guidelines in the development of *semantic* dialogue systems [12, 23]; over the last years, we have adhered strictly to the developed rule "No presentation without representation." The idea is to implement a generic, and semantic, dialogue shell that can be configured for and applied to domain-specific dia-

[1] http://www.w3.org/TR/voicexml20/
[2] http://www.w3.org/TR/emma/
[3] http://www.voicexml.org

logue applications. All messages transferred between internal and external components are based on RDF data structures which are modelled in a discourse ontology (also cf. [6, 8, 21]). Our systems for industrial dissemination have four main properties: (1) multimodality of user interaction, (2) ontological representation of interaction structures and queries, (3) semantic representation of the interface, and (4) encapsulation of the dialogue proper from the rest of the application.[4] These architectural decisions are based partly on usability issues that arise when dealing with end-to-end dialogue-based interaction systems for industrial dissemination (they correspond to the use case requirements). Intelligent AI systems that involve intelligent algorithms for dialogue processing and interaction management must be judged for their suitability in industrial environments. Our major concern which we observed in the development process for industrial applications over the last years is that the incorporation of AI technologies such as complex natural language understanding components (e.g., HPSG based speech understanding) and open-domain question answering functionality can unintentionally diminish a dialogue system's usability. This is because negative side-effects such as diminished predictability of what the system is doing at the moment and lost controllability of the internal dialogue processes (e.g., a question answering process) occur more often when AI components are involved. This tendency delivers new requirements for usability to account for the special demands introduced by the use of AI. For the identification of these usability issues, we adopted the binocular view of interactive intelligent systems (discussed in detail in [9]).

The main architectural challenges we encountered in implementing a new dialogue application for a new domain can be summarised as follows: first, providing a common basis for task-specific processing; second, accessing the entire application backend via a layered approach. In our experience, these challenges can best be solved by implementing the core of a dialogue runtime environment, an ontology dialogue platform (ODP) framework and its platform API (the DFKI spin-off company SemVox, see www.semvox.de, offers a commercial version), as well as providing configurable adaptor components. These translate between conventional answer data structures and ontology-based representations (in the case of, e.g., a SPARQL backend repository) or Web Services (WS)—ranging from simple HTTP-based REST services to Semantic Web Services, driven by declarative specifications [24]. Our ODP workbench (figure 1) builds upon the industry standard Eclipse and also integrates other established open source software development tools to support dialogue application development, automated testing, and interactive debugging. A distinguishing feature of the toolbox is the built-in support for eTFS (extended Typed Feature Structures), the optimised ODP-internal data representation for knowledge structures. This enables ontology-aware tools for the knowledge engineer and application developer. A detailed description of the rapid dialogue engineering process, which is possible thanks to the Eclipse plugins and templates, can be found in [25].

[4] A comprehensive overview of ontology-based dialogue processing and the systematic realisation of these properties can be found in [21], pp.71-131.

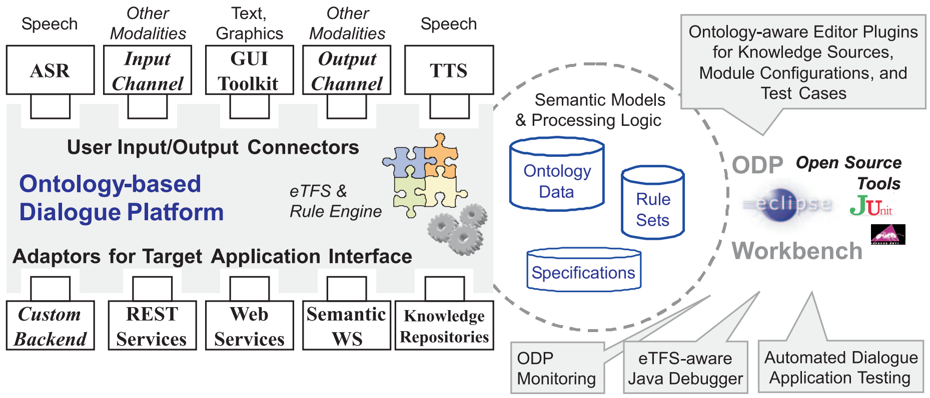

Fig. 1. Overall design of the discourse and dialogue infrastructure: ontology-based dialogue processing framework and workbench

2 Dialogue Tasks and Components

In addition to automatic speech recognition (ASR), dialogue tasks include the interpretation of the speech signal and other input modalities, the context-based generation of multimedia presentations, and the modelling of discourse structures. These topic areas are all part of the general research and development agenda within the area of discourse and dialogue with an emphasis on dialogue systems. According to the utility issues and industrial user requirements we identified (system robustness/usability and processing transparency play the major roles), we distinguished five dialogue tasks for these topics and built five dialogue components, respectively. These components which allow for a semantic and pragmatic/application-based modelling of discourse and dialogue are presented in more detail in the rest of this section.

Multimodal Semantic Processing and Navigation This task provides the rule-based fusion of different input modalities such as text, speech, and pointing gestures. We use a production-rules-based fusion and discourse engine which follows the implementation in [13]. Within the dialogue infrastructure, this component plays the major role since it provides basic and configurable dialogue processing capabilities that can be adapted to specific industrial application scenarios. More processing robustness is achieved through the application of a special robust parsing feature in the context of RDF graphs as a result of the input parsing process. When the user only utters catchwords instead of complete utterances, the semantic relationship between the catchwords can be guessed (following [26]) according to the ontological domain model of the industrial application domain. The Tool Suite builds upon the industry standard Eclipse and also integrates other established open source software development tools to support dialogue application development, automated testing, and interactive debugging.

The semantic navigation and interaction task/module builds the connection to backend services (the tasks Interactive Semantic Mediation and Web Service Composition) and the presentation module (the task Semantic Output representation) and allows for a graph-like presentation of incremental results (also cf. [2]). The spoken dialogue input is used to generate SPARQL queries on ontology instances (using a Sesame repository, see www.openrdf.org). Users are then presented a perspective on the result RDF graph with navigation possibilities.

Interactive Semantic Mediation Interactive semantic mediation has two aspects: (1) the mediation with the processing backend, the so-called dynamic knowledge base layer (i.e., the heterogeneous data repositories), and (2) the mediation with the user (advanced user interface functionality for the adaptation to new industrial use case scenarios). We developed a semantic mediation component to mediate between the query interpretation created in the dialogue shell and the semantic background services prevalent in the industry application context. The mediator can be used to, e.g., collect answers from external information sources. Since we deal with ontology-based information in heterogeneous terminology, *ontology matching* has become one of the major requirements. The ontology matching task can be addressed by several string-based or structure-based techniques (cf. [5], p.341, for example). As a new contribution in the context of large-scale speech-based AI systems, we think of ontology matching as a dialogue-based *interactive* mediation process (cognitive support frameworks for ontology mapping involve users). The overall approach is depicted in figure 2. A dialogue-based approach can make more use of partial, unsure mappings. Furthermore, it increases the usability in dialogue scenarios where the primary task is different from the matching task itself (cf. industrial usability requirements). So as not to annoy the user, he/she is presented only the difficult cases for disambiguation feedback; thus we use the dialogue shell to basically confirm or reject pre-considered alignments [20].

User Adaptation and Personalisation Industrial usability requirements clearly advocate new AI systems to work with user models and personalised information. An additional, related concern is given by the industrial requirement to provide user privacy protection and implement data security guidelines. This makes a user adaptation and personalisation task an indispensable asset for many advanced AI systems. Our discourse and dialogue infrastructure should benefit from user preferences, communities, their ontological modelling, the privacy issues of anonymity, data security, as well as data control by the user (also cf. processing transparency and transaction security for processes that involve sensitive data like medical patient records). Figure 3 illustrates the close interrelationships between the adaptation and personalisation blocks. In the context of new dialogue system infrastructures, persistent and personalised annotation of personal data (e.g., personal image annotations) should be provided. Further

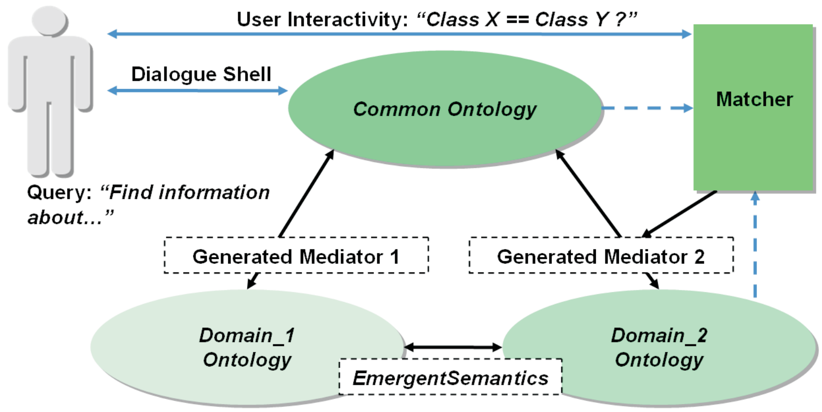

Fig. 2. Interactive semantic mediation for industry adaptation

topics of interest include interaction histories with dialogue session, task, and personalised interaction hierarchies. For example, our use case implementation tries to guess a predefined user group according to the interaction history.

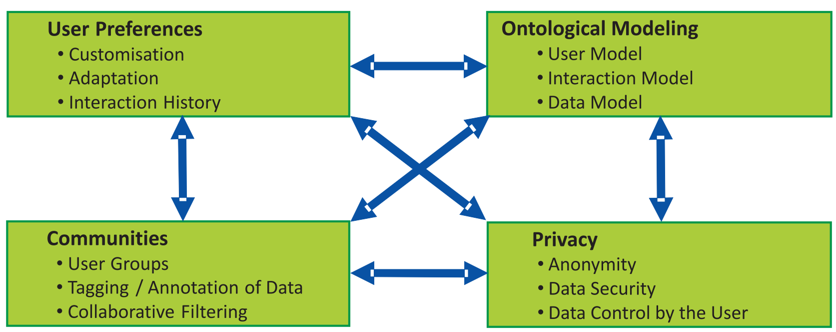

Fig. 3. Industry-relevant forms of user adaptation and personalisation

Interactive Service Composition The service composer module takes input from the multimodal semantic base module. A mapping of the structured representation of the query to a formal query is done first. The formal representation does not use concepts of verbs, nouns, etc. anymore, but rather uses the custom knowledge representation in RDF. If only one query answering backend exists, the knowledge representation of this backend can be used. Otherwise, an interim RDF-based representation is generated. The formal query is analysed and mapped to one or more services that can answer (parts of) the query. This step typically involves several substeps including decompositing the query into smaller parts, finding suitable services (service discovery), mapping the query to other representations, planning a series of service executions, and initiating possibly clarification requests or disambiguation requests. The different workflows form different routes in a pre-specified execution plan which includes the possibility to request clarification information from the user if needed (hence interactive service composition). Figure 4 outlines the hard-wired execution plan to dynamically address and compose SOAP/WSDL-based and REST-based services.

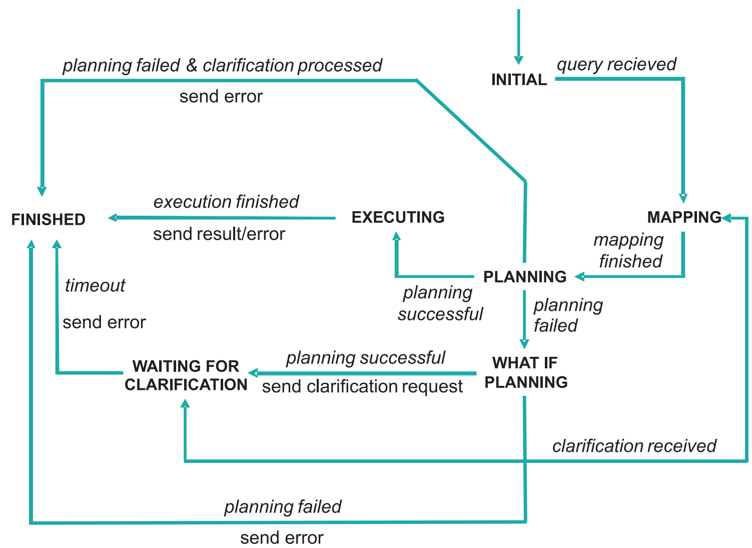

Fig. 4. Execution plan for (interactive) service composition

Semantic Output Representation We implemented a semantic output representation module which realises an abstract container concept called Semantic Interface Elements (SIEs) for the representation and activation of multimedia elements visible on, e.g., a touchscreen user interface [22]. The semantic output representation architecture comprises of several GUI-related submodules (such as the Semantic Interface Elements Manager or the Event Manager), dialogue-engine-related modules (such as the Interaction Manager or natural language Parser), and the Presentation Manager sub-module (i.e., the GUI Model Manager). The most important part of the architecture is the Display Manager which observes the behaviour of the currently displayed SIE. The display manager dispatches XML based messages to the dialogue system with the help of a message

Decoder/Encoder. This display manager has to be customised for every new application, whereas all other modules are generic. The Multimodal Dialogue Engine then processes the requests in the Interaction and Interpretation Manager modules. A new system action or reaction is then initiated and sent to the Presentation Manager. The GUI Model Manager builds up a presentation message in eTFS notation (internal typed feature structure format) before the complete message is sent to the Display Manager as a PreML message.

3 Industrial Dissemination

In the following discussion of industrial dissemination, the first example shows a medical prototype application for a radiologist while the second shows a business-to-business mobile application. Both dialogue systems build upon the multimodal speech-based discourse and dialogue infrastructure described in this paper. The prototypes put emphasis on various and combined input forms on different interaction devices, i.e., multitouch screens and iPhones.

Radiology Dialogue System In the MEDICO use case, we work on the direct industrial dissemination of a medical dialogue system prototype. Clinical care and research increasingly rely on digitised patient information. There is a growing need to store and organise all patient data, including health records, laboratory reports, and medical images. Effective retrieval of images builds on the semantic annotation of image contents. At the same time it is crucial that clinicians have access to a coherent view of these data within their particular diagnosis or treatment context. This means that with traditional user interfaces, users may browse or explore visualised patient data, but little or no help is given when it comes to the interpretation of what is being displayed. Semantic annotations should provide the necessary image information and a semantic dialogue shell should be used to ask questions about the image annotations and refine them while engaging the clinician in a natural speech dialogue at the same time.

The process of reading the images is highly efficient. Recently, structured reporting was introduced that allows radiologists to use predefined standardised forms for a limited but growing number of specific examinations. However, radiologists feel restricted by these standardised forms and fear a decrease in focus and eye dwell time on the images [7, 28]. As a result, the acceptance for structured reporting is still low among radiologists while referring physicians and hospital administrative staff are generally supportive of structured standardised reporting since it eases the communication with the radiologists and can be used more easily for further processing. We strive to overcome the limitations of structured reporting by allowing content-based information to be automatically extracted from medical images and (in combination with dialogue-based reporting) eliminating radiologist's requirements to fill out forms by enabling them to focus on the image while either dictating the image annotations of the reports to the dialogue system or refining existing annotations.

The domain-specific dialogue application [24], which uses a touchscreen (figure 5, left) to display the medical SIE windows, is able to process the following medical user-system dialogue:

1 **U:** "Show me the CTs, last examination, patient XY."
2 **S:** Shows corresponding patient CT studies as DICOM picture series and MR videos.
3 **U:** "Show me the internal organs: lungs, liver, then spleen and colon."
4 **S:** Shows corresponding patient image data according to referral record.
5 **U:** "This lymph node here (+ pointing gesture) is enlarged; so *lymphoblastic*. Are there any comparative cases in the hospital?"
6 **S:** "The search obtained this list of patients with similar lesions."
7 **U:** "Ah okay."
Our system switches to the comparative records to help the radiologist in the differential diagnosis of the suspicious case, before the next organ (liver) is examined.
8 **U:** "Find similar liver lesions with the characteristics: hyper-intense and/or coarse texture ..."
9 **S:** Our system again displays the search results ranked by the similarity and matching of the medical ontology terms that constrain the semantic search.

Fig. 5. Left: Multimodal touchscreen interface (reprinted from [24]). The clinician can touch the items and ask questions about them. Right: Mobile speech client for the business expert.

Currently, the prototype application is being tested in a clinical environment (University Hospitals Erlangen). Furthermore, the question of how to integrate this information and image knowledge with other types of data, such as patient data, is paramount. In a further step, individual, speech-based findings should be organised according to a specific body region and respective textual patient data.

Mobile Business-to-Business Dialogue System In the TEXO use case, we try to assist an employee and his superior in a production pipeline [15]. Our mobile business scenario is as follows: searching on a service platform, an employee of a company has found a suitable service which he needs for only a short period of time for his current work. Since he is not allowed to carry out the purchase, he formally requests the service by writing a ticket in the company-internal Enterprise Resource Planning (ERP) system. In the defined business process, only his superior can approve the request and buy the service. But first, the person in charge has to check for alternative services on the service platform which might be more suitable for the company in terms of quality or cost standards. The person in charge is currently away on business but he carries his mobile device with him that allows him to carry out the transaction on the go. The interaction is speech-based and employs a distributed version and instance of the dialogue system infrastructure. The mobile client (figure 5, left part of right half) streams the speech and click input to the dialogue server where the input fusion and reaction tasks are performed. The user can ask for specific services by saying, "Show me alternative services" or naming different service classes. After the ranked list of services is presented, multiple filter criteria can be specified at once (e.g., "Sort the results according to price and rating."). As a result, the services are displayed in a 2D grid which eases the selection according to multiple sorting criteria (figure 5, right).

We tested the mobile system in 12 business-related subtasks before the industrial dissemination. Eleven participants were recruited from a set of 50 people who responded to our request (only that fraction was found suitable). The selected people were all students (most of them in business or economics studies). From our analysis of the questionnaires we conclude that our mobile B2B system can be valuable for business users. Almost all users successfully completed the subtasks (89% of a total of 132 subtasks). In addition, many of them also provided the following positive feedback: they felt confident about the ticket purchase being successful. Then we introduced the prototype to the industrial partner. According to first user tests with the industrial partner SAP, the mobile speech client enhances the perceived usability of mobile business services and has a positive impact on mobile work productivity. More precisely, our mobile business application reflects the business task but simplifies the selection of services and the commitment of a transaction (internal purchase); it additionally minimises text entry (a limitation on mobile devices) and displays relevant information in such a way that a user is able to grasp it at first glance (3D visualisation).

4 Conclusion

Based on an integration platform for off-the-shelf dialogue solutions and internal dialogue modules, we described the parts of a new discourse and dialogue infrastructure. The ontology-based dialogue platform provides a technical solution for the dissemination challenge into industrial environments. The requirements for an industrial dissemination are met by implementing generic components for the most important tasks that we identified: multimodal semantic processing, semantic navigation, interactive semantic mediation, user adaptation/personalisation, interactive service composition, and semantic output representation.

Semantic (ontology-based) interpretations of dialogue utterances may become the key advancement in semantic search and dialogue-based interaction for industrial applications, thereby mediating and addressing dynamic, business-relevant information sources.

Acknowledgements Thanks go out to Robert Nesselrath, Yajing Zang, Markus Löckelt, Malte Kiesel, Matthieu Deru, Simon Bergweiler, Eugenie Giesbrecht, Alassane Ndiaye, Norbert Pfleger, Alexander Pfalzgraf, Jan Schehl, Jochen Steigner and Colette Weihrauch for the implementation and evaluation of the dialogue infrastructure. This research has been supported by the THESEUS Programme funded by the German Federal Ministry of Economics and Technology (01MQ07016).

References

1. J. Allen, D. Byron, M. Dzikovska, G. Ferguson, L. Galescu, and A. Stent. An Architecture for a Generic Dialogue Shell. *Natural Language Engineering*, 6(3):1–16, 2000.
2. R. Boselli and F. D. Paoli. Semantic navigation through multiple topic ontologies. In *Proceedings of Semantic Web Applications and Perspectives (SWAP)*, Trento, Italy, 2005.
3. P. Cimiano, P. Haase, and J. Heizmann. Porting natural language interfaces between domains: an experimental user study with the orakel system. In *IUI '07: Proceedings of the 12th international conference on Intelligent user interfaces*, pages 180–189, New York, NY, USA, 2007. ACM.
4. R. Engel. SPIN: A Semantic Parser for Spoken Dialog Systems. In *Proceedings of the 5th Slovenian and First International Language Technology Conference (IS-LTC 2006)*, 2006.
5. J. Euzenat and P. Shvaiko. *Ontology matching*. Springer-Verlag, Heidelberg, 2007.
6. D. Fensel, J. A. Hendler, H. Lieberman, and W. Wahlster, editors. *Spinning the Semantic Web: Bringing the World Wide Web to Its Full Potential*. MIT Press, 2003.
7. F. M. Hall. The radiology report of the future. *Radiology*, 251(2):313–316, 2009.
8. P. Hitzler, M. Krötzsch, and S. Rudolph. *Foundations of Semantic Web Technologies*. Chapman & Hall/CRC, August 2009.
9. A. D. Jameson, A. Spaulding, and N. Yorke-Smith. Introduction to the Special Issue on Usable AI. *AI Magazine*, 3(4):11–16, 2009.
10. V. Lopez, M. Pasin, and E. Motta. Aqualog: An ontology-portable question answering system for the semantic web. In *Proceedings of ESWC*, pages 546–562, 2005.
11. D. Martin, A. Cheyer, and D. Moran. The Open Agent Architecture: a framework for building distributed software systems. *Applied Artificial Intelligence*, 13(1/2):91–128, 1999.
12. S. Oviatt. Ten myths of multimodal interaction. *Communications of the ACM*, 42(11):74–81, 1999.
13. N. Pfleger. FADE - An Integrated Approach to Multimodal Fusion and Discourse Processing. In *Proceedings of the Dotoral Spotlight at ICMI 2005*, Trento, Italy, 2005.

14. R. Pieraccini and J. Huerta. Where do we go from here? research and commercial spoken dialog systems. In *Proceedings of the 6th SIGDdial Worlshop on Discourse and Dialogue*, pages 1–10, September 2005.
15. D. Porta, D. Sonntag, and R. Neßelrath. A Multimodal Mobile B2B Dialogue Interface on the iPhone. In *Proceedings of the 4th Workshop on Speech in Mobile and Pervasive Environments (SiMPE '09) in conjunction with MobileHCI '09*. ACM, 2009.
16. V. A. Ramachandran and I. Krishnamurthi. Nlion: Natural language interface for querying ontologies. In *COMPUTE '09: Proceedings of the 2nd Bangalore Annual Compute Conference on 2nd Bangalore Annual Compute Conference*, pages 1–4, New York, NY, USA, 2009. ACM.
17. N. Reithinger, D. Fedeler, A. Kumar, C. Lauer, E. Pecourt, and L. Romary. MIAMM - A Multimodal Dialogue System Using Haptics. In J. van Kuppevelt, L. Dybkjaer, and N. O. Bernsen, editors, *Advances in Natural Multimodal Dialogue Systems*. Springer, 2005.
18. N. Reithinger and D. Sonntag. An integration framework for a mobile multimodal dialogue system accessing the Semantic Web. In *Proceedings of INTERSPEECH*, pages 841–844, Lisbon, Portugal, 2005.
19. S. Seneff, R. Lau, and J. Polifroni. Organization, Communication, and Control in the Galaxy-II Conversational System. In *Proceedings of Eurospeech'99*, pages 1271–1274, Budapest, Hungary, 1999.
20. D. Sonntag. Towards dialogue-based interactive semantic mediation in the medical domain. In *Proceedings of the Third International Workshop on Ontology Matching (OM) collocated with the 7th International Semantic Web Conference (ISWC)*, 2008.
21. D. Sonntag. *Ontologies and Adaptivity in Dialogue for Question Answering*. AKA and IOS Press, Heidelberg, 2010.
22. D. Sonntag, M. Deru, and S. Bergweiler. Design and Implementation of Combined Mobile and Touchscreen-Based Multimodal Web 3.0 Interfaces. In *Proceedings of the International Conference on Artificial Intelligence (ICAI)*, pages 974–979, 2009.
23. D. Sonntag, R. Engel, G. Herzog, A. Pfalzgraf, N. Pfleger, M. Romanelli, and N. Reithinger. *SmartWeb Handheld—Multimodal Interaction with Ontological Knowledge Bases and Semantic Web Services*, volume 4451 of *Lecture Notes in Computer Science*, pages 272–295. Springer, 2007.
24. D. Sonntag and M. Möller. Unifying semantic annotation and querying in biomedical image repositories. In *Proceedings of International Conference on Knowledge Management and Information Sharing (KMIS)*, 2009.
25. D. Sonntag, G. Sonnenberg, R. Nesselrath, and G. Herzog. Supporting a rapid dialogue engineering process. In *Proceedings of the First International Workshop On Spoken Dialogue Systems Technology (IWSDS)*, 2009.
26. T. Tran, H. Wang, S. Rudolph, and P. Cimiano. Top-k Exploration of Query Candidates for Efficient Keyword Search on Graph-Shaped (RDF) Data. In *ICDE '09: Proceedings of the 2009 IEEE International Conference on Data Engineering*, pages 405–416, Washington, DC, USA, 2009. IEEE Computer Society.
27. W. Wahlster. SmartKom: Symmetric Multimodality in an Adaptive and Reusable Dialogue Shell. In R. Krahl and D. Günther, editors, *Proceedings of the Human Computer Interaction Status Conference 2003*, pages 47–62, Berlin, Germany, 2003. DLR.
28. D. L. Weiss and C. Langlotz. Structured reporting: Patient care enhancement or productivity nightmare? *Radiology*, 249(3):739–747, 2008.