A System for Facilitating and Enhancing Web Search

Steffen Staab[1], Christian Braun[2], Ilvio Bruder[3], Antje Düsterhöft[4],
Andreas Heuer[4], Meike Klettke[4], Günter Neumann[2], Bernd Prager[3],
Jan Pretzel[3], Hans-Peter Schnurr[1], Rudi Studer[1], Hans Uszkoreit[2], Burkhard Wrenger[3]

[1]AIFB, Univ. Karlsruhe, D-76128 Karlsruhe
[2]DFKI, Saarbrücken
[3]GECKO mbH, Rostock
[4]Informatik, Univ. Rostock

http://www.getess.gecko.de

Abstract

We present a system that uses semantic methods and natural language processing capabilites in order to provide comprehensive and easy-to-use access to tourist information in the WWW. Thereby, the system is designed such that as background knowledge and linguistic coverage increase, the benefits of the system improve, while it guarantees state-of-the-art information and database retrieval capabilities as its bottom line.

1 Introduction

Due to the vast amounts of information in the WWW, its users have more and more difficulties finding the information they are looking for among the many heterogeneous information resources. Therefore, methods for comfortable and intelligent access are in the primary focus of a number of research communities these days. Currently, syntactic methods of information retrieval prevail in realistic scenarios (cf., e.g., Ballerini et al. (1996)), such as in general search engines like AltaVista, but the limits inherent in these approaches often make finding the proper information a nuisance. On the other end of methodologies, semantic methods could provide just the right level for finding information, but they rely on explicitly annotated sources (cf., e.g., (Fensel et al., 1998)) or on complete and correct natural language understanding systems, both of which cannot be expected in the near future.

Therefore our system, GETESS, uses the semantics of documents in the WWW — as far as it is provided explicitly or as it can be inferred by an incomplete natural language understanding system, but relies on syntactic retrieval methods, once the methods at the semantic level fail to fulfill their task. In particular, we consider an information finding system that, *(i)*, has semantic knowledge for supporting the retrieval task, *(ii)*, partially, but robustly, understands natural language, *(iii)*, allows for several ways of interaction that appear natural to the human user, and, *(iv)*, combines knowledge from unstructured and semi-structured documents with knowledge from relational database systems.

In our project, we decided to aim at an information system that provides information finding and filtering methods for a restricted domain, *viz.* for prospective tourists that may travel in a certain region and are looking for all kinds of information, such as housing, leisure activities, seesights, etc. The information about all this cannot be found within a narrowly restricted format — neither in a single database nor in a single web site. Rather, the information agent must gather information that is stored on many different web servers, often in unstructured text, and even in some databases, such as a booking database of a hotel chain. In order to improve on common information retrieval systems, at least part of what is stated in the (HTML) texts must be made available semantically. However, since automatic text understanding is still far from perfect, we pursue a *fail-soft* approach that is based on extracting knowledge from text with a robust parser, but also integrates and falls back onto common information retrieval mechanisms when the more elaborate understanding component fails.

In the following, we draft the architecture of the GETESS-system with its overall sharing of the work load. From this outline we will then motivate and describe some key issues of the major subsystems of GETESS.

2 Architecture

The front end of the GETESS system (cf. a depiction of its architecture in Figure 1) provides a user interface that is embedded in a *dialogue system* controlling the history of interactions (cf. Section 3). Single interactions are handed to the query processor that selects the corresponding analysis methods, *viz.* the natural language processing module (NLP system; also cf. Section 5) or the information retrieval and database query mechanisms (cf. Section 4). While the latter ones can be directly used as input to the *search system*, the natural language processing module first translates the natural language query into a corresponding database query, before it sends this formal query to the search system.

In order to process queries and search for results, three kinds of resources are provided by the back end of the GETESS system. First, archived information is available in several content databases (the abstract DB, the index DB and the DB repository), the function of which is explained below. Second, the lexicon and the ontology provide metaknowledge about the queries, *viz.* about the grammatical status of words and their conceptual denotations. Third, a database incorporating dialogue sequences and user profiles, gives control over dialogue interactions.

While dialogue sequences and user profiles are acquired during the course of interactions and the metaknowledge is provided by the human modeller with the help of knowledge acquisition tools (*KA Tools*), the content databases must be filled automatically, since the contents of typical web sites change almost on a daily basis. For this task the *gatherer* searches regularly through relevant XML/HTML-pages and specified databases in order to generate corresponding entries in the abstract database, the index database and the database repository.

The content in the abstract database is derived from a robust, though incomplete natural language understanding module that parses documents and extracts semantic information building a so called "*abstract*" for each document. These abstracts are sets of facts, *i.e.* tupels, like `hasChurch(Alicante,Church-1)`, that could be extracted from natural language text, like "Alicante's major church was build during mediaval times." The index generator builds access information for full text search with information retrieval methods, while the DB repository

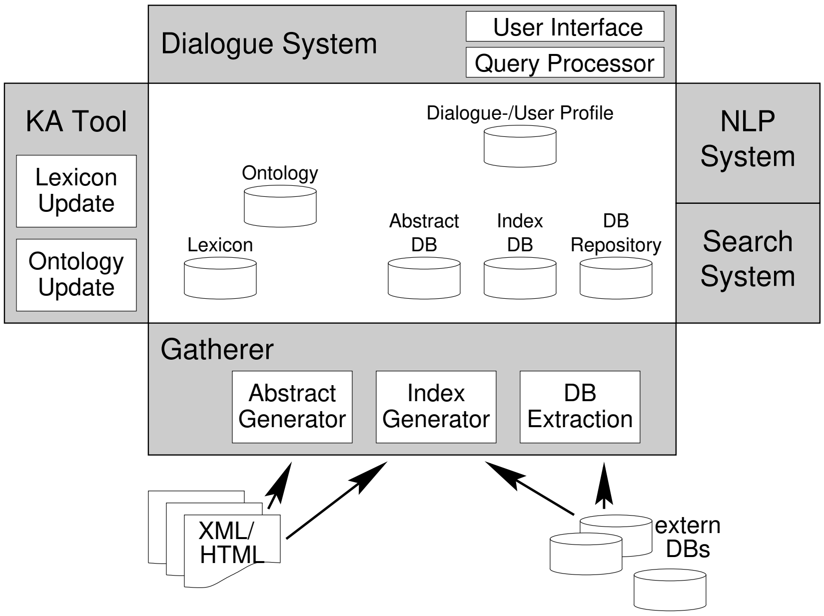

Figure 1: The GETESS system architecture

offers relevant views onto extern databases.

Subsequently, we will first introduce the front end, the dialogue system. The key issues here are concerned with *facilitating* user interaction at different levels of expertise (Section 3). At the back end of the system, the tools for gathering, database management and information retrieval provide the technical platform for efficiently updating and accessing the system's information repositories (Section 4). The natural language processing component in GETESS is employed by the dialogue system as well as by the back end in order to understand natural language queries and extract information from natural language texts, respectively, and, thus, *enhance* the quality of the web search (Section 5). Finally, we outline the function of the ontology that constitutes the "glue" of the system at the semantic level (Section 6).

3 Dialogue System

The dialogue system constitutes the interface between the human user and the data in the repositories of the GETESS system. In order to facilitate the user's task of finding the information he is looking for, users should be able to express queries conveniently at their level of expertise. This means the system should allow for intuitive interaction by natural language queries as well as for formal queries that may be the preferred mode of interaction by a human expert user. Indepent of the concrete mode of interaction, the system should react quickly and accurately while using the capabilities of the different modal actions.

In the GETESS system we allow for four types of interaction, *viz.* natural language, graphical interface, keyword search and formal database query. Since the methods for natural language processing as well as for keyword search and formal database queries form major components of the system, their description has been delegated to subsequent sections (5 as well as 4, resp.).

Thus, this chapter serves the following three goals: First, it is described how reasoning about user interactions may support the user's goal in quickly finding the appropriate information. Second, it is sketched how single interactions are treated as elements of a complex dialogue. Hence, the user does not have to start from scratch every time he inititiates a new query, but can instead refer to his previous queries, e.g. by requests like "Show me information related to this matter." where "*this matter*" relates to the last query. Finally, we give a glimpse on the use of the graphical interface.

The Knowledge Base of the Dialogue System. Knowledge is crucial for all modes of interaction, because we want the system to give appropriate responses to the user when problems arise. For instance, when a query results in an abundance of hits, the system must reason about why this problem might have occured and how it might be solved. Knowledge that allows for this type of reasoning is encoded in the *knowledge base of the dialogue system* (KBD).

The KBD includes all the definitions available in the ontology (cf. Section 6). These definitions help in explaining to the user why a query was too unspecific or giving him hints how he might try to rephrase the query such that he gets the information he is looking for. For example, if the user seeks for information on the local offers of "entertainment", the hit rate for a database query can be reduced by the choice of one of the refined search terms "music events", "theater events" and "sport events". Vice versa, a too specific choice like "folk music event" might result in no hits, but the hint towards more general search terms like "folk culture presentation" might bring up an event that also includes live demonstrations of "folk music". Further help is also provided through important terminological links such as synonyms, homonyms, antonyms and terms that may be parts of other terms, e.g. the show of a magician may be part of a circus show and, hence, the circus show might be a viable entertainment alternative to a magician's show.

In addition to the definitions of the ontology, the KBD features definitions and rules about dialogue concepts. At the moment, this part is tuned to map different interactions onto common requests to the database.[1] For example, the user input "I am looking for the station", which may also be supplemented by restrictions from the graphical interface, has the same meaning, i.e. it constitutes the same speech act, as "where is the station?". Therefore, both inputs must be mapped onto the same query to the database.

Both types of knowledge provide user support that reduces the number of inquiries the user has to pose to the system and, hence, accelerates the dialogue compared to common keyword retrieval interactions.

Complex Dialogues. As indicated above, information finding rarely produces an instantaneous hit — after the user has formulated just a single query. This is true for syntactic methods and it will improve only to a limited amount with semantic methods either. However, we believe that when a user's sequence of interactions is perceived as being executed in order to achieve a goal, then this task of finding the proper information can be substantially facilitated. For this purpose, we provide a query processor that analyses not only the single interactions, but also views them as being embedded into a more global structure.

The methodology we use is based on work done by Ahrenberg et al. (1996), who structure the dialog hierarchically into segments that are opened by a request and closed by the appropriate answer. The assumption in our scenario is that users typically have a request for a certain piece

[1]In the linguistic literature this mapping is defined by the way natural language propositions, requests or questions can be considered as so called *speech acts* (Austin, 1962).

of information and give related information in order to succeed. For example, they give a *topic*[2], which here boils down to a type restriction, like "sightseeing tour", and temporal information when they want to take part in a sightseeing tour during a particular time frame. The task of the dialogue system lies in zooming in or out on relevant information according to the interaction initiated by the user. For example, two user interactions[3] like, *(i)*, "Show me all theater events.", and, *(ii)*,"No, just the ones in August." return a large set of documents first (with feedback such as described in the previous subsection), but a much smaller set of data after the second interaction has narrowed down the focus.

Hence, we here identify *dialogue segments*, *interactions* and *topics* as the major parameters (though not the only ones) that the dialogue system keeps track of. By this way the user's single interactions may all convey towards the common information finding goal and, thus, facilitate the human-computer interaction.

The Graphical Interface. Besides of the natural language query capabilities and the possibility of directly composing a formal query, the GETESS system features a graphical interface. This interface constitutes an intermediate level of access to the system between the most professional (and fastest) one, *viz.* the formal query, and the most intuitive one (that requires a somewhat more elaborate interaction), *viz.* the natural language access. The graphical interface does not require from the user to learn the syntax of a particular query language or the concepts that are available in the ontology, but expects some basic understanding of formal systems from the user. This interface visualizes the ontology in a manner suited for selecting appropriate classes and attributes and, thus, allows the assembly of a formal query through simple mouse clicks. For this purpose, the ontology is visualized by a technology based on hyperbolic geometry (Lamping & Rao, 1996): classes in the center of the visualization are represented with a big circle, surrounding classes are represented with a smaller circle. This technique allows fast navigation to distant classes and a clear illustration of each class and its neighboring concepts.

4 Gathering, Database Management and Information Retrieval

In this section, we outline the back end of GETESS that gathers data from the web and stores it in a way that allows for efficient retrieval mechanisms as far as keyword search and formal queries are concerned.

The back end of GETESS employs a typical gatherer-broker structure, *viz.* a Harvest search service (Bowman et al., 1995) with a database interface. Though we use the tools provided by another project, SWING (Heyer et al., 1997), the setting of GETESS puts additional demands on the gatherer-broker system: *(i)*, the GETESS search engine has to work with facts contained in the abstracts, *(ii)*, ontology knowledge must be integrated into the process of analysing internet information as well as answering user queries, *(iii)*, internet information can be of different types (*e.g.*, HTML, XML texts), and, *(iv)*, data collections such as information stored in databases must also be accessible via the GETESS search engine. These different requirements must be

[2]The "topic" in a dialogue corresponds to what the dialogue is about. Usually, it is given only implicitly in natural language statements. In our setting it may also be given explicitly, *e.g.* through the graphical interface.

[3]Each interaction corresponds loosely to a speech act as introduced above, but might also be an act in the graphical user interface. Examples are *update* (users provide information to the system), *question*, *answer*, *assertion* or *directive*.

met both during the main process of gathering data and during the querying process (broker).

The Gatherer process. Periodically, internet information is analysed via internet agents in order to build a search index for the GETESS search engine. Information (*e.g.* HTML-texts, Postscript-files, ...) is checked to find keywords. Additionally, the GETESS-Gatherer has to build abstracts from these information. The two kinds of index data ('simple' keywords and abstracts) are stored in databases.

The Broker process. As indicated above, the dialogue system maps the user's queries (with the help of the natural language processing module and the definition in the ontology) onto formal or keyword queries in IRQL, the Information Retrieval Query Language. The IRQL language combines different kinds of queries — both database and information retrieval queries — thus providing access to the index data. The query result set is ranked with an user-centered ranking function. The ranked result set is then presented to the user via the dialogue system.

The integration of different types of information (full text, abstracts, relational database facts) during gathering and querying has put forth and still requires demand for research, however at the same time it raises new possibilities of posing queries, because:

1. Conventional search engines support an efficient search for keywords or combinations of keywords over a whole document. This is still possible, but the GETESS abstracts also relate information in the document to attributes. Exploiting this type of information, we can realise attributed queries. That means users have the possibility to search for terms in special attributes.

2. Searching for particular integer values, for instance prices or distances is nearly impossible with a conventional information retrieval approach. In GETESS, it will be possible to compare integer and real values, *e.g.* to search for all prices that fall below a threshold. In addition, one may also determine minimum, maximum and average values as well as sort and group results by particular values.

3. Database functionality brings up answers from the abstract databases that are composed of different abstracts. That means, for answering an query of a user we may refer and exploit facts derived from different websites. Thereby, it is not even necessary that these websites are connected by links, but all the algebraic operations given through database functionality can be employed in order to deduce information. For instance, today's cinema events may be announced in one web site, the corresponding reviews are found in another one. Database technology allows for retrieving all movies that are shown today and that received a good note in the corresponding review.

This functionality is implemented in an object relational database system. It will be employed in a distributed database solution, which provides for data storage at different local servers. Having made available different languages for accessing this repository of information, we are now researching a common language level, the IRQL described above, for accessing structured information (abstracts) and unstructured information (for instance HTML or XML) with the same interface. This language will then reduce the burden on the dialogue system, because the dialogue system will not have to distinguish between formal and keyword search anymore. Thus, IRQL will enhance the overall robustness of the system.

5 Natural Language Processing

In the GETESS system, the natural language processing (NLP) component is used in order to, *(i)*, linguistically analyse user queries specified in the dialogue system, *(ii)*, to generate the linguistic basis for the extraction of facts from NL-documents, and, *(iii)*, generate natural language responses from facts in the abstract database.

The design of this component is based on two major design criteria: First, the GETESS system requests a high degree of robustness and efficiency in order that the system may be applied in a real-world setting. The reason is that we must be able to process arbitrary sequences of strings efficiently, because "broken" documents appear on real web sites, and that the number of documents that must be processed is too large in order to allow for response times of several minutes per sentence. Second, we employ the same shallow NL core components and linguistic data managing tools for processing texts and extracting information as well as for analysing a user's query. Thus, we can reduce the amount of redundancies as far as possible and keep the system in a consistent state with regard to its language capabilities. For instance, when the internal linguistic representation of an NL query and the abstracts use the same data sources, and if we also use the same knowledge sources for NL-based generation, inconsistencies as a result of unshared data can be reduced.

For the purpose of a short presentation here, we abstract from two major components of the natural language processing component in GETESS. We do not elaborate on the natural language generation part that also includes features for summarizing facts from the abstract database. Moreover, we are well aware that our project serves an international tourist community and, therefore we will have to add multi-lingual access as well as multi-lingual presentations of the query results. However, at the current state of the project, we focus on parts of Germany as the touristic goal that we want to provide information about and, hence, focus on the analysis of German documents only.

Shallow text processing. The shallow text processor (STP) of GETESS is based on and an extension of SMES, an IE-core system developed at DFKI (see (Neumann et al., 1997; Neumann & Mazzini, 1998)). One of the major advantages of STP is that it makes a clean separation between domain independent and dependent knowledge sources. Its core components include: *(i)*, a text scanner, which recognizes, *e.g.*, number, date, time and word expressions, as well as text structure, like sentence markers and paragraphs; *(ii)*, a very fast component for morphological processing which performs inflectional analyses including processing of compounds and robust lexical access (e.g. analysing "houses" as the plural of "house"); *(iii)*, a chunk parser based on a cascade of weighted finite state transducers. The chunk parser performs recognition of phrases (specific phrases like complex date expressions and proper names, and general phrases, like nominal phrases and verb groups), collection of phrases into sentences, and determination of the grammatical functions (like the deep subject, which describes the acting person in "Tourist groups are led by a native guide.").

STP has large linguistic knowledge sources (e.g., 120.000 general stem entries, more than 20.000 verb–frame entries). The system is fast, and can process 400 words in about one second running all components. In order for adapting STP for GETESS, we have begun to evaluate STP's coverage on a corpus provided by our industrial partner. Though evaluation is blind, because the current knowledge sources have not been specified using any part of this corpus, we could analyse over 90% of all word forms and found that a majority of the remaining forms can be covered by domain specific lexica.

Extraction of facts. Finally, a word on the extraction of facts: The STP generates a linguistic analysis, *i.e.* it determines syntactic relations between words, *e.g.* between a verb and its subject. How these linguistic cues are exploited in order to go from natural language to a formal description is explained in the following section that elaborates on the semantic level of GETESS.

6 Ontology

As already mentioned, the gathering, use and querying of information with syntactical methods is very limited and in many cases not successful. A semantic reference model, an ontology, which structures the content and describes relationships between parts of the content helps to overcome these limitations. With the ontology in GETESS, we aim at two major purposes: First, it offers inference facilities that are exploited by the other modules, as, *e.g.*, described in Section 3, the dialogue module may ask for the types a particular instance belongs to in order to present alternative query options to the user. Second, the ontology acts as a mediator between the different modules. This latter role is explained here in more detail, since it illustrates how ontological design influences the working of the GETESS system and, in particular, the extraction of facts from natural language texts.

The text processing (cf. Section 5) of natural language documents and queries delivers syntactic relations between words and phrases. Whether and how this syntactic relation can be translated into a meaningful semantic relation, depends on how the tourism domain is conceptualized in the ontology. For example (cf. Fig. 2), the natural language processing system finds syntactic relations between the words "church" and "Alicante" in the phrase "Alicante's main church". The word "Alicante" refers to `Alicante` which is known as an instance of the class `city` in the database. The database refers to the ontology for the description of the classes `city` and `church`. Querying the ontology for semantic relations between `church` and `city` results in `hasBuilding` and `hasChurch`. Both relations are inherited from the class `location` to the class `city`. Since, `hasChurch` is the more specific one, a corresponding entry between `Alicante` and `church` is added to the abstract, *i.e.* the set of extracted facts, of the currently processed document in the abstract database.

This example shows that the design of the ontology determines the facts which may be extracted from texts, the database schema that must be used to store these facts and, thus of course, what information is made available at the semantic level. Hence, the ontology might constitute an engineering bottleneck. However, we try to overcome this problem by using the linguistic and statistical analyses of the text processing component for indicating frequent, though unmodelled, concepts and relations to the knowledge engineer.

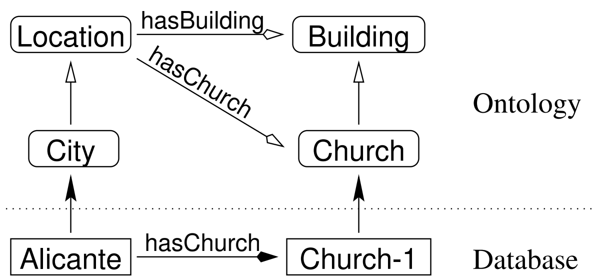

Figure 2: Interaction of Ontology with NLP system and Database Organization

7 Related Work

The GETESS project builds on and extends a lot of earlier work in various domains. In the natural language community, research like (Grosz et al., 1987; Wahlster et al., 1978) fostered the use of natural language application to databases, though these applications never reached the high preciseness and generality required in order to access typical databases, e.g. for accounting. Here, our approach seems better suited, since the imponderabilities of general natural language understanding are counterbalanced by information retrieval facilities and an accompanying graphical interface.

Only few researchers, e.g. Hahn et al. (1999), have elaborated on the interaction between natural language understanding and the corresponding use of ontologies. We think this to be an important point since underlying ontologies cannot only be used as submodules of text understanding systems, but can also be employed for a more direct access to the knowledge base and for providing an intermediate layer between text representation and extern databases, an interesting topic that has not been raised so far to the best of our knowledge.

As far as such queries of conceptual structures is concerned, we agree with McGuinness & Patel-Schneider (1998) that usability issues play a vital role in determining whether a semantic layer can be made available to the user and, hence, we elaborated on this topic early on (Fensel et al., 1998). We, thereby, keep in mind that regular users may find lengthy natural language questions too troublesome to deal with and, therefore, prefer an interface that allows fast access, but which is still more comfortable than any formal query language.

Projects that compare directly to GETESS are, e.g., Paradime (Neumann et al., 1997)[4], MULINEX (Capstick et al., 1998) and MIETTA (Buitelaar et al., 1998). However, none of these projects combines information extraction with similarly rich interactions at the semantic layer. Hence, to the best of our knowledge we are the only one integrating unstructured, semi-structured and highly-structured data with a variety of easy-to-use facilities for human-computer interaction.

8 Conclusion

In the project GETESS (GErman Text Exploitation and Search System) we decided to build an intelligent information finder that relies on current techniques for information retrieval and database querying as its bottom line. The support for finding informations is *enhanced* through an additional semantic layer that is based on ontological engineering and on a partial text understanding tool.

In order to *facilitate* web search, the dialogue is considered a complex entity. The analysis of sequences of interactions allows for refining, rephrasing or refocusing succeeding queries, and thus eliminates the burden of starting from scratch with every single interaction. Thereby, several modes of interaction are possible, besides keyword search and SQL-queries one can mix natural language queries with clicking in the graphical query interface.

Having built the single modules for our system, the next task in the GETESS project is bringing these components together. Given the design methodology of achieving entry level features

[4]Actually, GETESS uses the same linguistic core machinery as Paradime.

first and then working towards "the high ceiling" (*viz.* complete text understanding and representation), we expect benefits on the parts of economic and research interests early in the project. The system is general enough in order to be applied to many realistic scenarios, e.g. as an intelligent interface to a company's intranet, even though it is still far from offering a general solution for the most general information finding problems in the WWW. Further research will have to show an evaluation of how a user's performance in finding particular pieces of information looses or (hopefully) gains from using this information agent.

References

Ahrenberg, L., Dahlbäck, N., Jönnson, A., & Thure, A. (1996). Customizing interaction for natural language interfaces. *Computer and Information Science*, 1(1).

Austin, J. (1962). *How to Do Things with Words*. Oxford University Press.

Ballerini, J., Büchel, M., Knaus, D., Mateev, B., Mittendorf, M., Schäuble, P., Sheridan, P., & Wechsler, M. (1996). SPIDER retrieval system at TREC 5. In *Proc. of TREC-5. Gaithersburg, Maryland, November 20-22, 1996*. http://www-nlpir.nist.gov/pubs.

Bowman, C., Danzig, P., Hardy, R., Manber, U., & Schwartz, M. (1995). The harvest information discovery and access systems. *Networks and ISDN Systems*. http://ftp.cs.colorado.edu/pub/cs/techreports/schwartz/Harvest.Conf.ps.Z.

Buitelaar, P., Netter, K., & Xu, F. (1998). Integrating different strategies for cross-language information retrieval in the MIETTA project. In Hiemstra, D., de Jong, F., & Netter, K. (Eds.), *Language Technology in Multimedia Information Retrieval. Proceedings of the 14th Twente Workshop on Language Technology*, TWLT 14, pages 9–17. Universiteit Twente, Enschede.

Capstick, J., Diagne, A. K., Erbach, G., Uszkoreit, H., Cagno, F., Gadaleta, G., Hernandez, J. A., Korte, R., Leisenberg, A., Leisenberg, M., & Christ, O. (1998). MULINEX: Multilingual web search and navigation. In *Proceedings of Natural Language Processing and Industrial Applicatons*.

Fensel, D., Decker, S., Erdmann, M., & Studer, R. (1998). Ontobroker: The very high idea. In *FLAIRS-98: Proceedings of the 11th International Flairs Conference, Sanibal Island, Florida, May 1998*.

Grosz, B., Appelt, D., Martin, P., & Pereira, F. (1987). Team: An experiment in the design of transportable natural-language interfaces. *Artificial Intelligence*, 32(2):173–243.

Hahn, U., Romacker, M., & Schulz, S. (1999). How knowledge drives understanding: Matching medical ontologies with the needs of medical language processing. *AI in Medicine*, 15(1):25–51.

Heyer, A., Meyer, H., Düsterhöft, A., & Langer, U. (1997). SWING: Der Anfrage- und Suchdienst des Regionalen Informationssystems MV-Info. In *Tagungsband IuK-Tage Mecklenburg-Vorpommern. Schwerin, 27./28. Juni 1997*.

Lamping, J. & Rao, R. (1996). The hyperbolic browser: A focus + context technique for visualizing large hierarchies. *Journal of Visual Languages & Computing*, 7.

McGuinness, D. & Patel-Schneider, P. (1998). Usability issues in knowledge representation systems. In *Proc. of AAAI-98*, pages 608–614.

Neumann, G., Backofen, R., Baur, J., Becker, M., & Braun, C. (1997). An information extraction core system for real world german text processing. In *5th International Conference of Applied Natural Language Processing*, pages 208–215, Washington, USA.

Neumann, G. & Mazzini, G. (1998). Domain-adaptive information extraction. Technical report, DFKI, Saarbrücken.

Wahlster, W., Jameson, A., & Hoeppner, W. (1978). Glancing, referring and explaining in the dialogue system HAM-RPM. *American Journal of Computational Linguistics*, pages 53–67.