Natural Language Processing for Ambient Intelligence

Iryna Gurevych, Max Mühlhäuser

The goal of this contribution is to systematically investigate the possibilities of integrating concepts related to natural language processing (NLP) in Ambient Intelligence (AmI). Both research fields have been rapidly growing and evolving over recent years, substantially influencing the development of Computer Science as a whole. However, they are far from being sufficiently integrated yet. Utilizing NLP in AmI has the potential to generate cutting-edge research leading to substantial technological advances. In order to substantiate this claim, we review several application areas of NLP in AmI, such as service oriented computing, context aware systems, and natural human computer interfaces. This article introduces Ambient Intelligence as an exciting application for NLP researchers to stimulate explorative studies in this area. For AmI researchers, this article provides an insight into how NLP techniques can be employed to improve current Ambient Intelligence systems.

1 Introduction

An important area of artificial intelligence called **Natural Language Processing** deals with analyzing and purposefully manipulating human languages with the help of computers. NLP is closely related to Computational Linguistics, which employs computational tools to model phenomena in natural languages. NLP can be divided into two parts: **(i)** it develops core technologies for processing language in two forms, speech and text, to perform phonological, morphological, syntactic, semantic and pragmatic analysis, and **(ii)** it employs core technologies for processing text and speech, in order to build NLP applications, such as information retrieval, information extraction, question answering, summarization, machine translation, dialogue systems, etc.

As AmI systems are increasingly surrounded und penetrated by information, NLP techniques become highly relevant. For example, cell phones are used to access the Internet, digital cameras yield massive amounts of photos, smartphones can work with email, documents, etc. A substantial part of this information is represented in human languages (e.g., English, German) across a variety of media (e.g., text, sound, video) and is unstructured (e.g., textual service descriptions, news, or forum posts, product reviews, etc. in Web 2.0 information repositories). According to [1], *"the value of the software is proportional to the scale and dynamism of the data it helps to manage"*. Thus, we can conclude that the capability to process and make use of unstructured information can bring AmI systems to a much higher level of quality. Adding the respective functionality involves the ability to transform unstructured information in human languages into structured knowledge. This requires to determine the semantics of human expressions with the help of NLP and to make this knowledge available to AmI systems. Some studies have already begun to explore the exciting possibilities of integrating NLP and AmI technologies. Single papers highlighting individual research issues appear in the conferences about Ambient Intelligence, NLP, Semantic Web, Intelligent User Interfaces, and others [2]. Still, there is little exchange between these scientific communities and no established scientific forum with regular publications, as in other interdisciplinary areas [3].

Departing from this, our goal is **to identify and structure the relations** between NLP and AmI, resulting in a coherent view of interdisciplinary connections. As some of the readers may not be familiar with either of the two disciplines, we will first outline the scope and main research issues and directions in Ambient Intelligence and Natural Language Processing in Sections 2 and 3, respectively. After that, Section 4 will take a closer look at the relations existing between AmI and NLP and discuss the possibilities of using NLP in several areas of AmI. Some case studies from recent research literature will illustrate how NLP has so far been utilized in three AmI areas: **(i)** service oriented computing, **(ii)** context aware systems, and **(iii)** natural interfaces. We will conclude by summarizing the main benefits and limitations of NLP for Ambient Intelligence systems and outline possible developments in the future.

2 Ambient Intelligence

Ubiquitous Computing is *"a powerful shift in computation, where people live, work, and play in a seamlessly interweaving computing environment"* [4]. It is very close to **Ambient Intelligence**, which implies a seamless environment of computing, advanced networking technology and specific interfaces. This environment takes into account the specific characteristics of human users, considers their needs, and responds intelligently to spoken or multimodal interactions yielding intelligent dialogues. This technology is unobtrusive and often even invisible. Interacting with it should be enjoyable and intuitive for the user, avoiding the necessity of special training. Nowadays, AmI can no more be restricted to the areas of middleware, networking, and security. It goes beyond these areas and, therefore, requires the use of methods from artificial intelligence and human computer interaction to achieve high scalability, interoperability and ease of use for Ambient Intelligence systems.

AmI systems typically display many of the following properties. They are **embedded**, being an integral part of surrounding applications, which they control. They are **mobile**, being part of the moving application, e.g., a car, a bicycle, or a mobile device. They are **distributed**, which means that components intelligently cooperate with each other by exchanging useful information and services. Recently, service oriented architectures became widely adopted as an architectural paradigm, including issues, such as automatic service discovery, service selection and service composition. AmI systems make use of **contextual information**, such as temperature, pressure, objects, or people in the proximity, etc. This information is often acquired by appropriate sensors, which may constitute intelligent sensor networks. AmI systems are subject to **continuous update**, as it is necessary to keep the internal information up-to-date in order to fulfil their mission. Content based publish/subscribe systems is a popular implementation paradigm. *Adaptive* AmI systems should be equipped with **intelligent user interfaces**, which are customized regarding the application scenario and the end device employed in human computer interaction. E.g., speech may be an appropriate mode of interaction for mobile systems and complex ambient environments. AmI systems have to organize their services under the condition that the amount of information available varies and external services may not always be available. They are **heterogeneous**, i.e. they need to encompass a wide range of devices, such as cameras, scanners, printers, various home appliances, environmental sensors, specialized components for speech, handwriting, etc.

In summary, Ambient Intelligence systems have to address several **major problem domains** under one umbrella: **(i)** solutions for huge networks and applications and support for global spontaneous interoperability of software components, **(ii)** intelligent ad-hoc cooperation of a multitude of (specialized) devices over unreliable wireless networks, **(iii)** adaptivity and context awareness, **(iv)** problems in IT security, such as privacy and traceability, trust models, or legal binding of service users and providers, and **(v)** human computer interaction for AmI systems as a separate concern. NLP techniques are relevant to many of these problem domains, as it is described later in Section 4.

3 Natural Language Processing

3.1 Core Technologies

Several distinct NLP layers have to work together to enable automatic language analysis. For spoken language, the lowest layer performs phonological analysis, whereby the structure of a word is modelled as a sequence of phonemes. In morphological analysis, the word is represented as a sequence of morphemes, whereby a sentence is analyzed in terms of its syntactic constituents during the syntactic analysis. Semantic analysis assigns semantic interpretation to the constituents of the sentence, whereas the pragmatic analysis determines the communicative purpose of discourse units. A comprehensive overview and a detailed description of core NLP technologies is given in [5], which serves as a textbook in numerous NLP courses.

In the last decade, tremendous progress has been made in creating robust NLP tools for text analysis. Corpus driven approaches to language analysis became dominating in NLP. This was facilitated by large scale annotation initiatives and a wide adoption of machine learning techniques for modelling language phenomena. Part-of-speech taggers, e.g., TreeTagger [6], syntactic parsers, e.g., LOPAR [7], and information extraction tools [8] are now available. Furthermore, broad coverage lexical semantic resources were created, e.g., CELEX [9], or GermaNet [10]. They provide detailed, broad coverage linguistic information required to process unrestricted discourse. Semantic parsing [12] and discourse parsing [13, 14] are active areas of NLP research.

In order to build complex NLP systems and to ensure the reusability of the developed components, special middleware has been created. One particular example is the open-source framework Unstructured Information Management Architecture (UIMA) by IBM [15]. UIMA was created to build a bridge from unstructured information to structured knowledge. Similar to alternative NLP software architectures, such as GATE [16], UIMA provides central services, such as data storage, component communication and visualization of results. Additionally, a set of component integration and convenient implementation routines are provided. As UIMA is an industrial strength and scalable integration platform, it represents a good choice for composing NLP applications and integrating them in Ambient Intelligence systems. UIMA Component Repository [17] and Darmstadt Knowledge Processing Software Repository [18] contain off-the-shelf components, which can be employed for that.

3.2 NLP Applications

Major NLP applications relevant to Ambient Intelligence can be grouped into three categories: **(i)** semantic analysis (ontology learning, reasoning and mapping, text mining), **(ii)** information management (information retrieval, automatic summarization, information extraction, question answering), and **(iii)** speech processing and dialogue systems. Due to space limitations, we cannot cover all NLP technologies and exclude, for example, machine translation, handwriting recognition, or spelling correction. Furthermore, multiple NLP technologies can be combined to more complex applications, e.g., a multilingual question answering system (information retrieval, question answering, and machine translation), or speech-based information retrieval on mobile devices (speech recognition, information retrieval, automatic summarization, and text-to-speech).

Semantic analysis is the task of assigning semantic structure to unstructured information. One of the research areas in semantic analysis is *ontology learning*. Learning ontologies from the text entails automatic extraction of semantic concepts and taxonomic and other relations existing between them. For example, given a corpus of textual Web service descriptions, an ontology describing these services can be automatically induced. Typically, ontology learning involves machine learning methods and features based on language analysis, such as syntactic functions. Once an ontology has been created, reasoning can be done by inferencing over the ontology. For example, two different ontologies can be *aligned* with each other to establish correspondences between equivalent semantic concepts. The field of *text mining* is related to ontology learning. Text mining uses data mining techniques applied to natural language to discover previously unknown relationships.

Language based **information management** is the task of organizing unstructured language information in such a way, that it can be optimally prepared and presented to users according to their information needs. *Information retrieval* is a research area concerned with searching for relevant documents represented in different types of media and contained, e.g., in the Internet, relational stand-alone databases, or any other information repositories. Information retrieval should be tightly coupled with semantic processing, providing more flexible matching strategies than string based techniques. *Information extraction* is concerned with analyzing documents, detecting and extracting relevant pieces of information, such as the topic, participants, events, and their effects. *Automatic summarization* is a technology designed to create a short version of the language input preserving its meaning. In *question answering*, natural language queries are used to retrieve documents, detect the most relevant segments and possibly directly extract an answer to the original query. Thereby, information from possibly heterogeneous sources has to be found and aggregated to satisfy a user's information need.

Speech and dialogue processing encompass a set of topics related to analyzing spoken language and enabling natural forms of human computer interaction based on dialogue. *Speech recognition* is a technology to analyze spoken language input and transform it into an orthographical representation, which is further processed by an understanding component to derive an internal semantic representation of the input. *Text-to-speech* denotes the opposite mode of interaction, whereby the spoken output is automatically generated from an internal or an orthographic representation. *Natural language generation* is a research area concerned with generating natural language discourse from a semantic representation. For example, natural language instructions can be generated for a tourist from the underlying route representations. Finally, *dialogue systems* allow to carry out a dialogue with the user, i.e., to meaningfully process sequences of utterances pursuing some communicative dialogue goal, e.g., planning a trip or conducting a ticket reservation. The context of interaction including the dialogue history has to be accounted for to enable more natural interaction.

4 NLP for Ambient Intelligence

NLP techniques can be effectively utilized in many Ambient Intelligence tasks, since large-scale AmI faces the following challenges: **(i)** large amounts of information in textual form, which cannot be processed without automatic support; **(ii)** the necessity to integrate information from heterogeneous sources, and **(iii)** speech as the preferable mode of interaction in hands-free ambient and mobile environments and on small devices. Therefore, NLP adds great value to unstructured information by interpretation, transformation, filtering and augmentation of the data with semantic annotations.

Ambient Intelligence involves multiple research areas, as shown in Figure 1. NLP techniques can be utilized to enhance the capabilities of AmI systems by exploiting textual and speech information attached to them. In *service oriented computing*, meta data about services can be derived from their textual descriptions to achieve better discovery, scalability and interoperability. *Context aware systems* can exploit application related information sources in text or speech to dynamically update their context models with the aim of improved adaptivity. *Natural human computer interfaces* allow natural language based forms of communication, utilizing a broad range of speech and text technologies. In *peer-to-peer networks*, semantic information retrieval techniques are required to improve advanced information search and integration. *Event based systems* can adopt NLP techniques to develop more flexible and, therefore, scalable semantic techniques for message routing. Finally, *trust computing* can greatly benefit from analyzing community based information repositories in Web 2.0 and integrating the information extracted from user communities into trust models. Several case studies, where NLP is utilized in AmI will be presented below.

Similarly, AmI technologies can be effectively utilized in building NLP systems, as detailed in Figure 1, though a deep analysis of this is outside the scope of this paper. For example, AmI techniques can improve the architecture of NLP systems, or provide additional knowledge, such as sensed contextual information, to resolve natural language ambiguities. Therefore, there is a high potential for mutual benefit between both research fields. In the following, we discuss some case studies from recent research, where NLP has been applied to AmI in three areas: **(i)** service oriented computing; **(ii)** context aware systems, and **(iii)** natural human computer interaction.

4.1 Service Oriented Computing

Service oriented architectures (SOA) imply the use of loosely coupled services to build more complex software systems, in order to react to dynamic needs and compose re-usable and configurable services on demand. are **loosely coupled** and **interchangeable**. Therefore, a great degree of interoperability between them is desirable. Currently, this is addressed at the operational or syntactic level by the standards independent of the underlying platform or programming language (such as Java, .NET), e.g., SOAP (Simple Object Access Protocol), WSDL (Web Services Definition Language), or UDDI (Universal Description and Directory Service). Neither of these standards addresses the semantic level of interoperability, which is expected to play an increasingly important role due to a continuously growing number of services. Text classification and categorization, ontology aligning and information retrieval techniques can be employed for this purpose, as described below.

A machine learning approach utilizing iterative relational classification algorithm to learn the semantic category of a Web service is presented in [21]. They cast this task as a text classification problem, whereby the category is predicted based on the strings representing operations and datatypes. Given semantically annotated services as training data, the system learns to predict semantic labels for unseen instances of services. Despite the imperfect performance of the classification algorithm, the system can significantly reduce the effort of manual semantic annotation by providing useful suggestions to human experts. The algorithm is evaluated on 164 Web services. The software, experimental results and the annotated data for these experiments gathered from the repositories SALCentral.org and XMethods.com are available from the project homepage [22].

A framework for semi automatically annotating Web service descriptions with ontologies is described in [23]. The authors present an algorithm to match and annotate WSDL files with relevant domain ontologies. In the first step, XML schemata and ontologies are converted into a uniform graph representation called SchemaGraph. In the second step, every concept from the WSDL graph is compared with all concepts in the ontology graph with the help of a matching function composed of two different measures: an Element Level Match and Schema Level Match. The first one provides the linguistic similarity of two concepts, while the second one considers the structural similarity by evaluating the structure of the trees attached to them.

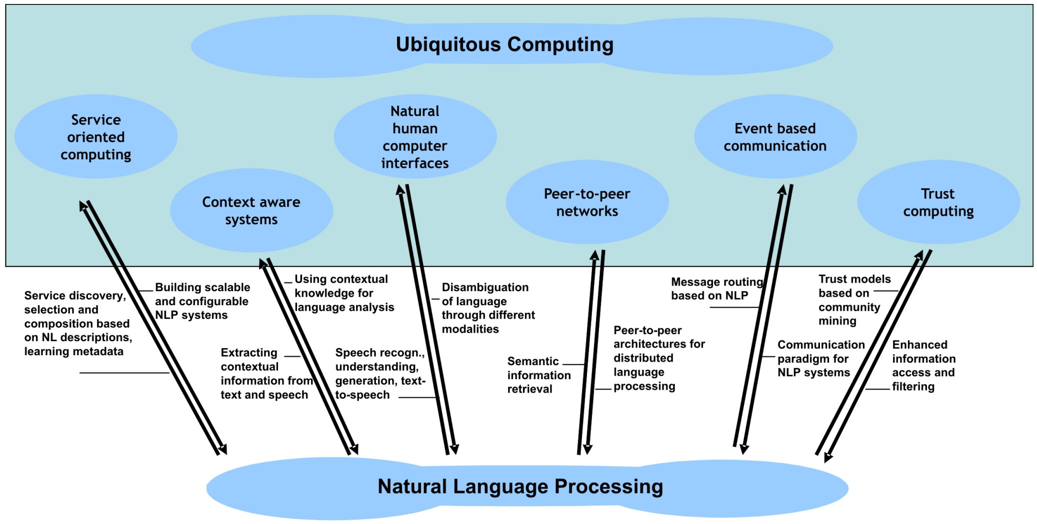

Figure 1: Interconnections between NLP and AmI. In the box, different areas of Ambient Intelligence are shown. The arrows show connections between AmI areas and NLP, with the labels describing example uses of either NLP or AmI technologies.

In the next step, each Web service description is compared with available ontologies and a set of mappings is created. The first one is called the Average Concept Match, and the second one is the Average Service Match. Similarly to the previously mentioned study [21], the authors present an empirical evaluation based on 424 Web services from SALCentral.org and XMethods.com. However, their approach incorporates more knowledge as it considers the structure of WSDL concepts, rather than just the names. Also, it uses ontologies for classification as opposed to vocabularies used in [21]. The disadvantage of this approach is its dependence on the availability of domain ontologies.

The problem of learning a semantic description of the service can be regarded as the task of ontology learning in NLP terms. Yet, generating semantic meta data is only part of the task. It is also necessary to make services interoperable, similar to the task of machine translation for people speaking different languages. The descriptions of services, which may be based on different ontologies, must be inter-related. This task can be understood as the problem of ontology mapping or aligning, which is a special kind of reasoning, whereby inferences are made about semantic relations existing between concepts in different ontologies. Most typically, concepts in the source ontology are translated into concepts in the target ontology. A number of surveys comprehensively describe the state-of-the-art in ontology aligning and the algorithms employed in this field [26, 27].

Given a specific task to be performed, discovery of Web services and their dynamic integration requires advanced searching capabilities for people and machines. The authors of [28] review current retrieval methods in Web service repositories, such as UDDI, Bindingpoint, .NET XML Web Services Repertory and others, and conclude that they **(i)** employ only simple keyword and substring search techniques leading to many irrelevant hits, **(ii)** offer limited browsing capabilities for human experts, as only few and low quality meta data is available, and **(iii)** the meta data is not fully exploited for presentation. Furthermore, it is necessary to consider not only the domain of activity, but also the functionality, the type of input and output and the restrictions that may apply to Web services.

Advanced search capabilities utilizing NLP techniques and semantic knowledge from broad coverage resources can provide great benefits for service providers, service brokers and service clients. Textual information attached to Web services, such as their names, descriptions, names of operations, etc. can be analyzed with NLP methods. In particular, a large body of research on semantically enhanced information retrieval techniques exist, which can be applied to the domain of Web service search [19, 20]. Searching for Web services based on textual information will also eliminate the need for service descriptions, based on manually specified ontologies that cannot easily be changed, e.g., if a service is modified. Further benefits from NLP arise in the area of service discovery. Thereby, services can be searched and found beyond one specific platform or community of users. This leads to new business opportunities and re-purposing of services for service marketplaces. In the future, such marketplaces are likely to integrate an increasing number of business intelligence elements. NLP-based analysis of communities, such as analyzing the trends and opinions about the usage of a specific service, can be utilized to improve service quality management and adapt service selection according to the opinions of users.

4.2 Context Aware Systems

Modelling and use of context in Ambient Intelligence applications is crucial to such fundamental tasks as creating explanations tailored to users, the acquisition of knowledge in its context of use, and the learning capability of AmI systems as part of their functionality. Typical sources of contextual information in AmI are sensorial input, databases, or an explicit model provided by the system designer. Deriving structured knowledge and inferences from the sensorial data is often non-trivial. Additionally, a lot of information is encoded in natural language as the most convenient form of human communication, e.g., users are alerted about road works or traffic jams in radio reports, or they can read about the place and the time of some interesting events in a newspaper. Therefore, natural language, either written or spoken, can provide an additional source of contextual information for AmI applications. For example, the study in [29] proposed to utilize the information derived from human language to close the gap between immediate sensorial updates and delayed manual updates of context models by using NLP techniques. The authors introduce the idea of a text sensor, which we use to explain several challenges to be accounted for, when extracting contextual information from natural language.

Challenge (i) Mapping free text to concepts in the context model often requires inference, as one entity may be described by different words, which is especially relevant for proper names, called named entities. For example, the *pope Benedict XVI* is sometimes referred to using his secular name *Joseph Ratzinger*, various combinations like *Kardinal Joseph Ratzinger* or *Joseph Kardinal Ratzinger*, and common misspellings like *Josef Ratzinger*. Additionally, abbreviations have to be resolved. Such tasks can be addressed by using NLP resources, such as lexicons listing the possible variations, and morphological components for translating the words to their base form. To compute the similarity of two strings, distance measures can be employed. The implementation of many distance measures is available from the open-source library SecondString [30]. Furthermore, Wikipedia has been recently recognized as an excellent source of background knowledge to support NLP [31]. With lexical-semantic and world knowledge extracted from Wikipedia (e.g., Wikipedia redirect pages encoding the information about synonymy, spelling variations and abbreviations), inferences can be performed. To access the knowledge encoded in Wikipedia, a high-performance Java-based API has been developed [32].

A further difficulty in mapping words to the context model is the resolution of deictic expressions and of temporal and spatial references. As opposed to named entities, e.g., *Darmstadt*, and definite descriptions, e.g., *the city*, deictic expressions refer to the personal, temporal, or spatial aspect of an utterance, e.g., *this part*, *here*. Therefore, their resolution heavily depends on the contextual knowledge. The same is true for resolving temporal and spatial references, such as *at the end of the year*, or *in the central part of the city*. To resolve such expressions with a high level of accuracy, domain and world knowledge represented in an ontology is typically required.

Challenge (ii) Representing uncertainty in contextual information, which can arise, e.g., due to possible errors in sensor measurements, outdated information, etc. The evidence obtained from different information sources may even be contradictory, e.g., a calendar entry may report a different location as the one determined with the help of a GPS component. As previously mentioned, ambiguity also constitutes the major problem in analyzing natural languages. E.g., the word *Sydney* can mean the name of the city or the name of the person. If it is used in the meaning of city, there is a problem of disambiguating which of the multiple instances of Sydney is meant, i.e., in Australia, in Japan, in Canada, in the Unites States, etc. For disambiguating the meaning of words, NLP developed specialized technologies, such as named entity recognition, or word sense disambiguation [5]. However, the information extracted by these components is subject to uncertainty, too, and has to be represented in the context model. This can be done by employing approaches proposed to represent uncertainty in the context model [33, 34].

Challenge (iii) Constructing hybrid models to accommodate multiple types of information, such as geometric, topological and ontological representations, and appropriately linking them to each other. This is important to provide a mapping between linguistic and geographic entities and to express the relations that exist between them. Also, NLP often needs access to the context model in order to improve the quality of the analysis. This is different to regular sensors, which do not make use of the knowledge in the context model, but only contribute to it.

Traffic Analyzer, the system presented in [29], detects traffic jams in web news feeds, uses this information to relate the extracted information to existing data, updates the context model and displays the information on the map. Its context model is managed by the NEXUS platform [35]. The analysis of WWW documents about traffic jams is performed using NLP techniques, such as information extraction. Thereby, a template is filled with relevant information, backed up with domain dependent application knowledge to support the analysis process. In the next step, the results are linked with the application model to detect context data, such as traffic jams, as geographical objects. TrafficAnalyzer is an example of how natural language and NLP techniques can be used in context aware AmI systems. The performance of NLP can be further improved by making the analysis components more robust and providing better knowledge sources.

4.3 Natural Human Computer Interfaces

AmI systems differ from conventional desktop computing application scenarios. They are often embedded in small devices used in a mobile mode or are distributed among multiple components in ambient environments. Therefore, traditional modes of interaction are often impossible. Speech and natural language interfaces become attractive for creating human computer interfaces in Ambient Intelligence. In recent years, many speech based and multimodal systems have been created. Speech based systems range from command-style to more complex spoken dialogue systems. Command-style interfaces are more appropriate in limited application domains, such as device control, while spoken dialogue systems are suitable to accomplish more complex tasks, such as information seeking [36, 37]. Major technological issues in designing speech interfaces are: (i) reducing recognition errors, and (ii) dealing with ambiguous input. Recognition errors can be minimized by combining different types of evidence, such as acoustic and linguistic features derived from the dialogue. Some works incorporated semantic information to evaluate speech recognition hypotheses, such as a slot language model [41] and ontology-based semantic coherence metrics [42]. Such approaches typically require rich knowledge in the target domain, which makes them less scalable.

Multimodal systems [39] support user interactions with the system, involving multiple channels of communication (also called modalities), such as voice, gesture, handwriting, text, etc. The main motivation behind designing multimodal AmI systems is to improve their accessibility on the input side and their usability on the output side. Natural and multimodal interaction can mean different things to various groups of users in different situations. It can mean portable access to multimedia, news and entertainment services with the convenience of speech for text entry from everywhere. In mobile contexts, e.g., in the car, it can mean an integrated dashboard system offering hands-free navigation, entertainment, news, and communications applications.

In homes, multimodal interaction may be applied to remote control of the home entertainment centre's integrated multimedia player or recorder for television, radio, music, video and games in a living room. In the office, finally, the user can choose how to interact with the computer, e.g., using a pen, keyboard, or spoken commands. Often, different modalities are used to disambiguate the meaning of the user's input. For example, gestures can be used to disambiguate the meaning of the deictic expression uttered by the user, e.g., *Give me that* [40].

Dealing with ambiguous input in the context of mobile and ambient environments requires methods to resolve deictic expressions. Information from multiple modalities can be employed to disambiguate speech. For example, a dynamic Bayesian Network is used to integrate the modalities in [43]. The system, called XWand, is able to interpret commands, such as *Switch on* issued by the user, while she is pointing at the lamp. Another approach uses physical context information collected with the help of sensors to support the disambiguation process. [44] presents a system, where the user points a PDA with a camera at an object tagged with a colour bar code. The bar code is recognized and this information is used to improve the results of the speech recognition. [45] utilizes information, such as tracking the user location, recognizing pointing by the finger or laser pointer at the screen, or reacting if a user sits on a particular piece of furniture in MIT's Intelligent Room. This information is compared with device states and interaction history to dynamically adjust speech recognition knowledge sources, such as grammars and lexicons, to improve the speech recognition accuracy. A similar approach is taken in [46], whereby different types of contextual information such as sensed (physical) context and discourse context are combined to make inferences about a meeting taking place in the room.

Synergistic combinations of NLP technologies, e.g., using speech recognition with information retrieval, or machine translation with speech recognition and question answering, would yield a multitude of innovative intelligent human computer interfaces.

5 Conclusions and Outlook

In this article, we argued that applying Natural Language Processing in Ambient Intelligence opens up a lot of exciting research issues. Recent advances in NLP make it feasible to integrate this technology in real-life AmI systems, whenever it is beneficial to exploit sources of textual or speech information for various purposes, e.g., data integration, semantic analysis, context modelling and natural human computer interfaces. NLP is an indispensable tool to cope with large amounts of information, by providing automated support to humans or machines, who would otherwise not be able to analyze the information.

Lately, a lot of research carried out in NLP was primarily focused on common evaluation benchmarks, leaving less room for *exploratory interdisciplinary research*. However, the transfer of established NLP methods to new domains represents a research issue in itself, due to special properties of the language involved. Lack of corpora for evaluations in these new domains makes a comparison of different NLP approaches difficult. The corpora employed in such evaluations are often created by the researchers themselves, and, therefore, suffer from methodological limitations. In order to achieve better integration of research in NLP and AmI, it is necessary to establish *a research community* on this topic. This will foster systematic exchange of the most recent research results between the two communities, to discover new interesting application areas and to establish more principled and objective empirical evaluations. We believe that using NLP for AmI is far from being well exploited and has an exciting future ahead of it.

References

[1] T. O'Reilly. What Is Web 2.0: Design Patterns and Business Models for the Next Generation of Software. http://www.oreillynet.com/pub/a/oreilly/tim/news/2005/09/30/what-is-web-20.html

[2] Pervasive Computing Conference http://www.percom.org/, Computational Linguistics Conference http://www.aclweb.org/, International Semantic Web Conference http://iswc.semanticweb.org/, International Conference on Intelligent User Interfaces http://www.iuiconf.org/

[3] International Artificial Intelligence in Education Society http://aied.inf.ed.ac.uk/aiedsoc.html

[4] M. Weiser. The Computer for the Twenty-First Century. Scientific American, pp. 94-10, September 1991.

[5] D. Jurafsky and J. H. Martin. SPEECH and LANGUAGE PROCESSING. An Introduction to Natural Language Processing, Computational Linguistics, and Speech Recognition. Second Edition. http://www.cs.colorado.edu/ martin/slp2.html

[6] H. Schmid. Probabilistic Part-of-Speech Tagging Using Decision Trees. In Daniel Jones and Harold Somers (Eds.) New Methods in Language Processing. Studies in Computational Linguistics, UCL Press, London, GB, pp. 154-164, 1997.

[7] H. Schmid. Efficient Parsing of Highly Ambiguous Context-Free Grammars with Bit Vectors. In Proc. of the 20th Int. Conf. on Computational Linguistics (COLING 2004). Geneva, Switzerland, 2004.

[8] G. Neumann and J. Piskorski. A Shallow Text Processing Core Engine. Journal of Computational Intelligence 18(3), 2002.

[9] http://www.ru.nl/celex/

[10] http://www.sfs.nphil.uni-tuebingen.de/lsd/

[11] Shared tasks and common evaluation benchmarks in NLP, s. http://www-nlpir.nist.gov/related_projects/muc/, http://www-nlpir.nist.gov/projects/duc/, http://www.senseval.org/, http://www.pascal-network.org/Challenges/RTE/Introduction/

[12] S. Pradhan, W. Ward, K. Hacioglu, J. Martin and D. Jurafsky. Shallow Semantic Parsing Using Support Vector Machines. In Proc. of the Human Language Technology Conf./North American chapter of the Association of Computational Linguistics (HLT/NAACL). Boston, MA, 2004.

[13] D. Marcu and A. Echihabi. An Unsupervised Approach to Recognizing Discourse Relations. In Proc. of the 40th Annual Meeting of the Association for Computational Linguistics (ACL-02). Philadelphia, PA, July 2002.

[14] K. Forbes, E. Miltsakaki, R. Prasad, A. Sarkar, A. Joshi and B. Webber. D-LTAG System: Discourse Parsing with a Lexicalized Tree-Adjoining Grammar. Journal of Logic, Language and Information 12(3). Kluwer Academic Publishers, Hingham, MA, USA, 2003, pp. 261-279.

[15] D. Ferrucci and A. Lally. Building an example application with the Unstructured Information Management Architecture. IBM Systems Journal: Unstructured Information Management 43(3), 2004, pp. 455-475.

[16] H. Cunningham, D. Maynard, K. Bontcheva and V. Tablan. GATE: A framework and graphical development environment for robust NLP tools and applications. In Proc. of the 40th Annual Meeting of the Association for Computational Linguistics (ACL-02). Philadelphia, PA, July 2002.

[17] http://uima.lti.cs.cmu.edu/

[18] http://www.ukp.tu-darmstadt.de/

[19] I. Gurevych. Anwendungen des semantischen Wissens über Konzepte im Information Retrieval. In Proc. of Knowledge eXtended: Die Kooperation von Wissenschaftlern, Bibliothekaren und IT-Spezialisten. Jülich, Germany, 2005.

[20] C. Müller and I. Gurevych. Exploring the Potential of Semantic Relatedness in Information Retrieval. In Proc. of LWA 2006 Lernen - Wissensentdeckung - Adaptivität: Information Retrieval. Hildesheim, Germany, 2006, pp. 126-131.

[21] A. Heß, E. Johnston, and N. Kushmerick. ASSAM: A Tool for Semi-Automatically Annotating Semantic Web Services. In 3rd Int. Semantic Web Conf. (ISWC 2004), 2004.

[22] http://www.few.vu.nl/ andreas/projects/annotator/

[23] A. Patil, S. Oundhakar, A. Sheth, K. Verma. METEOR-S Web service Annotation Framework. In Proc. of the World Wide Web Conf., July 2004.

[24] http://oaei.ontologymatching.org/

[25] http://www.ontologymatching.org/

[26] P. Shvaiko and J. Euzenat. A Survey of Schema-based Matching Approaches. Journal on Data Semantics, 2005.

[27] N. Noy. Semantic Integration: A Survey of Ontology-based Approaches. Sigmod Record, Special Issue on Semantic Integration, 2004.

[28] M. Sabou and J. Pan. Towards Improving Web Service Repositories through Semantic Web Techniques. Workshop on Semantic Web Enabled Software Engineering (SWESE) at ISWC'2005.

[29] A. Blessing, S. Klatt, D. Nicklas, S. Volz and H. Schütze. Language-Derived Information and Context Models. In Proc. of 3rd IEEE PerCom Workshop on Context Modeling and Reasoning (CoMoRea) at 4th IEEE Int. Conf. on Pervasive Computing and Communication (PerCom'06), January 2006.

[30] http://secondstring.sourceforge.net/

[31] T. Zesch, I. Gurevych and M. Mühlhäuser. Analyzing and Accessing Wikipedia as a Lexical Semantic Resource. In Biannual Conf. of the Society for Computational Linguistics and Language Technology (GLDV'2007), 2007.

[32] http://www.ukp.tu-darmstadt.de/software/WikipediaAPI

[33] A. Ranganathan, R. H. Campbell, A. Ravi and A. Mahajan. ConChat: A Context-Aware Chat Program. IEEE Pervasive Computing, 1(3), July 2002, p.p. 51-57.

[34] N. Hönle, U. Kappeler, D. Nicklas, T. Schwarz and M. Grossmann. Benefits of Integrating Meta Data into a Context Model. In Proc. of the Third IEEE Int. Conf. on Pervasive Computing and Communications Workshops, 2005, p.p. 25-29.

[35] D. Klinec and S. Volz. NEXUS - Positioning and Communication Environment for Spatially Aware Applications. In: ISPRS Congress 2000, IAPRS, Amsterdam, Netherlands, 2000, p.p. 324-330.

[36] V. Zue. Conversational Interfaces: Advances and Challenges. In Proc. of Eurospeech'97. Rhodes, Greece, 1997.

[37] S. Seneff. Response planning and generation in the Mercury flight reservation system. Computer Speech and Language 16, 2002, p.p. 283-312.

[39] W. Wahlster. Smartkom: Foundations of Multimodal Dialogue Systems. Springer, Berlin, 2006.

[40] S. Oviatt. Breaking the Robustness Barrier: Recent Progress on the Design of Robust Multimodal Systems. Advances in Computers (ed. by M. Zelkowitz). Academic Press, 2002, vol. 56, p.p. 305-341.

[41] S. Pradhan and W. Ward. Estimating Semantic Confidence for Spoken Dialogue Systems. In Proc. of Int. Conf. on Acoustic Speech and Signal Processing (ICASSP-2002). Orlando, Florida, May 2002.

[42] I. Gurevych and R. Porzel. Using knowledge-based scores for identifying best speech recognition hypothesis. In Proc. of ISCA Tutorial and Research Workshop on Error Handling in Spoken Dialogue Systems. Chateau-d'Oex-Vaud, Switzerland, 2003, p.p. 77-81.

[43] A. Wilson and S. Shafer. XWand: UI for Intelligent Spaces. In Proc. of SIGCHI, 2003.

[44] J. Rekimoto and K. Nagao. The World through the Computer: Computer Augmented Interaction with Real World Environments. In Proc. of UIST'95, 1995, p.p. 29-36.

[45] M. Coen, L. Weisman, K. Thomas, M. Groh and K. Gajos. A Context Sensitive Natural Language Modality for an Intelligent Room. In Proc. of the 1st Int. Workshop on Managing Interactions in Smart Environments (MANSE'99). Dublin, Ireland, 1999, p.p. 201-212.

[46] L. H. Leong, S. Kobayashi, N. Koshizuka and K. Sakamura. CASIS: a context-aware speech interface system. In Proc. of the 10th Int. Conf. on Intelligent User Interfaces. San Diego, California, USA, January 10-13, 2005.

Contact

Dr. Iryna Gurevych, Ubiquitous Knowledge Processing Group
Prof. Dr. Max Mühlhäuser, Telecooperation Division
Computer Science Department
Darmstadt University of Technology
Hochschulstr. 10, 64289 Darmstadt
Tel.: +49 (0)6151-165411
Fax: +49 (0)6151-163052
Email: {gurevych,max}@tk.informatik.tu-darmstadt.de

Bild

Iryna Gurevych is Head of the Ubiquitous Knowledge Processing Group at Darmstadt University of Technology. She has a PhD in Computational Linguistics and post-doc experience in the Mobile Assistance Systems and Natural Language Processing groups at EML Research (Heidelberg, Germany). She is principle investigator in several research projects in NLP, information retrieval and text mining, which focus on semantic information processing.

Bild

Max Mühlhäuser is Head of the Telecooperation Division at Darmstadt University of Technology. He initiated and is principal investigator of a large number of research projects in the area of Ambient Intelligence. In 1993, he founded one of the first and leading Ubiquitous Computing Labs in Europe, the TeCO institute (www.teco.edu) in Karlsruhe, Germany. He has published more than 200 articles.