elsnews 12.4

The Newsletter of the European Network in Human Language Technologies Winter 2003/4

Towards a Roadmap for Speech Technology

José M. Pardo, Universidad Politécnica Madrid

A special session at Eurospeech 2003 in Geneva, September 4 2003

This special session at Eurospeech 2003, organised by ELSNET, was intended to contribute to the wider roadmapping exercise reported previously in *ELSNews* The session consisted of an introduction to the Roadmap concept and the objective of the session by Steven Krauwer, followed by four invited papers presented by Paul Heisterkamp, Bjom Granstrom, Ron Cole, and Roger Moore. The session ended with a discussion of the topics presented and of the general roadmap exercise, with questions from the audience.

Invited papers

The paper by **Heisterkamp** addressed some of the problems and solutions that we encounter today in spoken dialogue systems (SDS). He mentioned how we should teach people to use the systems. Most of the actual problems are due to people not speaking to the system, people not saying what they mean, and people not providing the information requested by the system. Heisterkamp gave some examples where semantic comprehension would be very difficult to achieve by SDS for a long time to come, many more than ten years. People don't always mean what they say, neither do they always say what they mean.

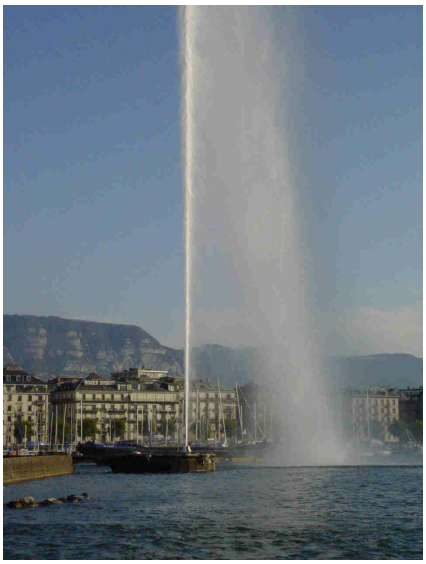

Geneva's famous fountain, symbol of Eurospeech 2003

Heistercamp also addressed the fact that many systems we use today work successfully because of conventions, not necessarily through logical and natural behaviour. For instance the 'qwerty' layout of the typewriter was not only designed to help typists write faster, but to avoid mechanic conflicts between consecutive keys. Today it is a convention and everybody uses it.

One of the conclusions is that it would be good to establish SDS conventions instead of trying to make a machine mimic human behaviour exactly. The ultimate goal of a SDS should be that of a better, cheaper, convenient, and reliable service, instead of trying to match the human process. Naturalness and ease of use are not necessarily the same and to know how to use a system we need conventions and training, not necessarily a natural system. He also addressed the important point of investing in dialogue design, a topic not always taken into consideration, that contributes to the failure of many systems.

The paper by **Granstrom** covered some of the problems related to the use of multimodality today (i.e., integrating audio and visual modalities) and how to solve

Note to our readers

With the transition to Framework Programme 6, future sponsorship of ELSNET, and hence the future of ELSNews, was still in the balance as this issue went to press. We are working hard to try to ensure that ELSNews continues its service to language and speech researchers in Europe and beyond, as it has now for twelve years. Up-to-date news will be placed on the ELSNET web site, www.elsnet.org, as it becomes available.

Inside this issue:

Winter 2003/4

elsnet

ISSN 1350-990X

Participants enjoy exploring Lake Geneva (picture courtesy of Dan Bohus)

ELSNews is published at the Department of Informatics, University of Sussex. It is printed by the University of Sussex Print Unit.

Editor: Lynne Cahill.

Editorial Team: Geoffrey Sampson, Steven Krauwer, and Brigitte Burger.

ISSN 1350-990X

FOR INFORMATION
Contributions to *ELSNews*, and corrections, should be sent to:
lynneca@sussex.ac.uk

Tel: +44 1273 678659
Fax: +44 1273 671320

Material for the next issue is due:
March 19th 2004

them. The paper discussed some of the possible models for using multiple signals and integrating them in a complete communication process. By its nature, the paper presented the state-of-the-art of actual systems (it did not attempt to predict the future). Concretely it presented three problems: how to obtain data, how to model them, and how to exploit them in dialogue systems. In the presentation, some demonstration of facial synthesis was done, emphasising the holistic nature of the speech communication process. Three applications of facial synthesis were presented.

The paper by **Cole** set out his concept of the roadmap. First, we have to decide on the objective and next on the kind of journey we want to make to arrive at the objective. It correctly, in my opinion, set the objective: to achieve 'Great Communication'; and the kind of journey: characterised by much exploration, and guided by successes and failures during these adventures, led by ambitious goals and conducted by independent researchers. Under this point of view it would be impossible to establish predictions because they depend very much on the successes and failures of the researchers, among other parameters. The paper also points out the parameters that define 'Great Communication': emotional, immersive, and personal. It hypothesised also that the evaluation of future systems, taking this view into account, would be related more to the usefulness or not of the experience of the users using the systems.

The paper reminded us about the multidisciplinarity of the problem: speech research, psychology and cognitive sciences, linguistics, computer science, and electrical engineering. It is important to establish a good interdisciplinary team with experts in all these disciplines. It also showed some steps that the Colorado team are taking in this direction. Cole's opinion clearly envisages future systems closer and closer to human behaviour.

The paper by **Moore** was dedicated to a particular roadmap exercise on spoken language output. It had all the ingredients that we are looking for: we want to know what will happen in the mid-term future in the area and the possible steps needed to make it happen.

In contrast with Cole's vision, the Roadmap is defined by the objective (where to go?) and the optimum way to achieve the objective (how to get there?) and not the nature of the trip (what kind of journey we want to get there?). It was also driven by market pull, trying to match it to technology push. This view is much more practical and realistic about what will possibly happen and it also matches any reasonable plan for an industry involved in the field. First the market opportunities are identified, then the product feature concepts that could satisfy them are defined, and finally the technical solutions required to realise the new products.

The market drivers are identified from the 6th EU framework programme: "a future in which computers and networks will be integrated into the everyday environment, rendering accessible a multitude of services and applications through easy-to use human interfaces". Although not time-labelled, some technical

challenges are listed in the spoken language output task: improved modelling of style, voice, and prosody, better modelling of the vocal tract, and – very interesting – models that learn, models with proprioceptive feedback that hear and monitor their own performance.

Other papers presented at Eurospeech relevant to the Roadmap exercise

At least three other papers at Eurospeech 2003 were relevant to the roadmap exercise:

"Speech and Language Processing: Where Have We Been and Where Are We Going?" *(Kenneth Ward Church)*

This paper speculated about the future with regard to several questions, in particular:
1. Is more data better data? The progress of language processing has been alternating between data models and knowledge models. At the start we had data models (twenty years ago) then there was a move towards knowledge models (grammars, rules) to constrain the data models. We are currently in an era of data models again (now in the second decade of it). Church's prediction is that in ten years time we will have to go back to knowledge-based models.
2. What will we do with the petabytes (10^{15}) of data that will be available? Search will become a key problem and models that address this issue will be important.

"ISCA Special Session: Hot Topics in Speech Synthesis" *(Gerard Bailly, Nick Campbell, Bernd Möbius)*

What are the hot topics for speech synthesis? How will they differ in five years time? ISCA's Special Interest Group on Synthesis (SynSIG) presents a few suggestions. This paper attempted to identify the top five hot topics, based not on an analysis of what is being presented at current workshops and conferences, but rather on an analysis of what is not. It was accompanied by results from a questionnaire polling SynSIG members' views and opinions. The fact that it abstracts the opinions of several recognized experts in the area makes it meaningful. It mentions evaluation, extension, emotion, multimodality, and "type of input to the synthesiser" as key topics today that will still be alive in 2008.

"A Comparison of the Data Requirements of Automatic Speech Recognition Systems and Human Listeners" *(Roger K. Moore)*

This was another contribution from Roger Moore on an important topic, this time speculating about the speech recognition task.

Since the introduction of hidden Markov modelling there has been an increasing emphasis on data-driven approaches to automatic speech recognition. This derives from the fact that systems trained on substantial corpora readily outperform those that rely on more phonetic or linguistic priors. Similarly, extra training data almost always results in a reduction in word error rate – "there's no data like more data". Moore asked (and answered) the following questions:
How much speech does a human listen to in a lifetime? *120,000 hours*. How much speech would be needed with actual system performance (extrapolating) to achieve human performance? *70 lifetimes*. His conclusion was that our model of speech recognition training is much poorer than a humans, so more work would be need to be done to improve language models.


José M. Pardo

Discussion

Some of the discussion was related to the presentations on multimodal animated agents. One question that arose from the audience was whether the goal is to strive for naturalness of the agent, believability, or simply a system that can help us in any way, even if it is not similar to a human being. The answer is that society will assess what applications are possible. Another comment on the topic was that the animated agent will change the way a human person speaks to it, so they could finally sustain effective communication.

Roger Moore warned us about the term naturalness. Naturalness is good but it can be a serious problem. He warns against using this term, as it is unclear what is considered natural. Human beings are constantly evolving and what is natural today might not have been natural yesterday.

I shall finish with Ron Cole's conclusion on his vision of the roadmap: "Take big challenges and solve them".

FOR INFORMATION

José M. Pardo is Head of the Speech Technology Group at Universidad Politécnica Madrid

Email: pardo@die.upm.es

Web: www.die.upm.es/personal/J-pardo.html.en

More information about Eurospeech can be found at www.symporg.ch/eurospeech/

More information about the roadmap can be found on the ELSNET web site: www.elsnet.org

Opinion

Language Engineering One

Annie Zaenen, *Xerox PARC*

Language engineering is a term that has been in fashion since the mid-nineties. It is supposed to be distinguished from Computational Linguistics, which is more theoretical, and Natural Language Processing, which covers both theory and applications. Language Engineering in general is, as Hamish Cunningham[1] put it, "the discipline or act of engineering software systems that perform tasks involving processing human language. Both the construction process and its outputs are measurable and predictable. The literature of the field relates to both application of relevant scientific results and a body of practice."

The term is, however, also used in a more restricted way as in the title of the recent *Handbook for Language Engineers*[2], which addresses mainly those involved in linguistic data analysis within the broader enterprise.

To avoid confusion, I will call the work done within LE that concentrates on the analysis and representation of specific characteristics of natural language 'LE_l'. This includes rule writing of various sorts, annotation of training corpora, designing features for statistical systems, etc. Although about ten years ago linguistic knowledge was seen as an obstacle to progress by some, in fact it has been appealed to constantly even in statistical applications. In some applications its contribution is of such a low level that it doesn't require linguistic training. For instance, every college-educated person knows what parts of speech are and for certain applications it might be good enough to get by with those given in a dictionary. As NLP applications move away from basic information retrieval towards more tasks that require some natural language understanding, a Language Engineer$_l$ faces more interesting challenges. For some of these tasks a degree in linguistics might be a good starting point, e.g., syntactic rule writing or syntactic treebank encoding of corpora. But for other essential tasks most linguists are not well prepared, nor is anybody else. Let's take as an example task 'co-reference resolution'. This is recognised as an essential task for a great number of NL applications and was some years ago the object of a couple of MUC competitions. As far as I can see, the LE_l done in that context was dismal. The problems with it are gently documented in a couple of papers by Kibble and van Deemter[3]. To give just one of the most striking examples, the annotators were told that the relation that they had to annotate for was transitive, but also that it held between the 'the stock price' and '$55' in a sentence such as "The stock price reached $55." To their credit, the developers of the scheme noticed that this would lead to problems in the case of sentences such as "The stock price fell from $4.02 to $3.85" and added the fix that in such cases the annotators should choose the last number. As van Deemter and Kibble observe, what if the next sentence is something like "Later that day it fell to an even lower value, at $3.82"?

It wouldn't be worth talking about this, of course, if it were an isolated incident. I have been looking at several annotation schemes and the output of several LE_l endeavours lately, and in most cases the experience is rather depressing. Is this because LE_lers are stupid (e.g., because the only linguists that do this work are the ones that can't find a real linguistics job)? I don't think so. I know LE_lers that are superb linguists, and I am sure that the developers of the co-reference annotation guidelines saw that they were painting themselves into a corner. But they most likely had a deadline. This is the first of the two main reasons that I see for the sorry state of some of LE_l. It has to do with the economics of the enterprise; NL applications have to be cheap and even when a national standard enterprise such as NIST gets into the act, the credo is that annotations have to be done quickly and cheaply, never mind that it will cost more to cope with the results of this cheap work than to do it right in the first place. The second is that more and more of the information that one tries to get to through these annotations (or rules, or features) has no science behind it because it has not been of concern to traditional linguists. Only very few linguists will finish their degree (of whatever level it may be) with a clear knowledge of the different types of anaphoric relations that can hold between an anaphor and its antecedent, and the way these anaphoric relations determine or don't determine co-reference or whichever other relations one would like to establish between entities. One of the reasons is that most of syntax and semantics has been concentrating on sentence-internal relations, another reason is that it has been concentrating on

phenomena that are assumed to be categorical. (There are academic subfields that in the best of all worlds could provide some theoretical underpinnings for some of the problems LE_1 is struggling with, but they are at this point not pursued in a computational setting.)

This gap in the study of natural language phenomena makes the situation of LE_1 particularly precarious. Yes, LE_1ers have to learn about corpus building, ontologies, evaluation, etc. – all the useful topics covered in the *Handbook for Language Engineers* – but that doesn't fill the scientific gap. One might not need deep 'explanatory adequacy'. To a point, Engineering can get by without scientific explanation, but it cannot get by without science that is descriptively adequate. As long as this problem is not recognised and taken care of, the Language Engineer is on a mission impossible, and LE in the broad sense on the way to new disappointments.

FOR INFORMATION

Annie Zaenen is Principal Research Scientist at Xerox PARC, USA
Email: zaenen@parc.com
Web: www2.parc.com/istl/members/zaenen

References:

[1]Cunningham, H. (1999) A Definition and Short History of Language Engineering. *Journal of Natural Language Engineering*, 5(1):1-16.
[2]Ali Farghaly, ed. (2003) *Handbook for Language Engineers*, CSLI Publications.
[3] E.g. van Deemter, Kees and Rodger Kibble (2000) On Coreferring: Coreference in MUC and Related Annotation Schemes, *Computational Linguistics* 26:4; 629-38

Announcement

Journal of Phonetics Special Issue

Temporal Integration in the Perception of Speech

Edited by Sarah Hawkins and Noël Nguyen

Temporal integration in speech perception refers to how chunks of information passed from the ear to the brain at different times are linked together by a listener to decode meanings. Until recently, this entailed the perceptual grouping of acoustic cues to identify phonetic segments. However, in recent years speech perception has come to be studied within a much broader context and with a much more inter- and multidisciplinary perspective.

Based on the TIPS workshop held in Aix-en-Provence in April 2002, this volume is much more than a collection of papers. Thirteen of the fifteen papers presented at the workshop are included (the other two being published elsewhere), along with no fewer than fourteen commentary papers comparing and relating the content of the focus papers. The two papers not included here are discussed in the commentary papers.

What makes this volume particularly interesting is the interdisciplinary nature of much of the work, with psycholinguistics, psychoacoustics, neuropsychology, and computational modelling covered in addition to the more traditional subject matter of the *Journal of Phonetics*. Papers that fall outside the core subject matter include tutorial material, so that they are comprehensible to those unfamiliar with these other fields.

The volume is dedicated to the memory of Peter Jusczyk, a pioneer in the field who tragically died before the workshop, at which he was due to give an invited talk.

Contributors:
Robert R. Remez; Stephen D. Goldinger and Tamiko Azuma; John Local; Gerard J. Docherty; Lynne C. Nygaard; John Coleman; Sarah Hawkins; Simon King; John Laver; Sophie K. Scott; Stephen Grossberg; M. Gareth Gaskell; Steven Greenberg, Hannah Carvey, Leah Hitchcock, and Shuangyu Chang; David W. Gow Jr; Shihab Shamma; Betty Tuller; Stuart Rosen; Maria Mody; Nathalie Bedoin; Alain de Cheveigné; Brian C. J. Moore; Françoise Macar; Martin Cooke; Anne Christophe, Ariel Gout, Sharon Peperkamp, and James Morgan; Robert F. Port; Catherine T. Best; Christopher T. Kello.

FOR INFORMATION

The special issue is available online at www.sciencedirect.com/science/journal/00954470

More information can be found at the Elsevier web site: wwwelsevier.com/locate/phonetics

The abstracts for the original workshop papers can be found online at: www.lpl.univ-aix.fr/~tips and are also available in a book published by the University of Cambridge Printing Service (ISBN: 1680-8908)

Collaborative Interface Agents from 2D to 3D

Candace Sidner, *MERL*

At Mitsubishi Electric Research Labs (MERL), a small team of researchers has been developing interface agents that can collaborate conversationally based on existing computational linguistic theory[1]. This team has developed Collagen™, a Java-based collaborative agent middleware tool, that has been used in more than ten different domain applications. These range from agent assistance to a human user, where human decision making dominates the interaction, to tutoring a human student [3], where agent decision-making dominates. Collaborators and users of Collagen include USC-ISI, Technical University of Delft, MIT Media Lab, and Mitre Corporation. The Collagen effort shares goals and conceptual frameworks in common with many other dialogue-based research efforts, e.g., the WITAS work of Lemon and Peters.

Middleware tools make it possible to build systems, but as with many other tools (such as TRINDIKIT), the challenge lies in building applications with that tool. The MERL group takes the approach of developing new applications with Collagen in order to understand what aspects of a dialogue system to grow next. Recent applications are focused on multimodal (speech and touch) mixed-initiative web form filling (see web site below), and assistance to customers using living environment products such as personal video recorders[4] and new generation thermostats[1].

These efforts over the past nine years have involved agents based in the 2D world of the interface screen. But recently, collaboration in the 3D world has become a focus of the group. For 3D, the agent is a robot with the ability to gesture with its arms, body and head. This effort shares some concerns in common with work on embodied conversational agents. The theoretical challenge is to understand the engagement process in conversation and collaboration: how do collaborators use both language and gesture to indicate their connection and their attention to one another? And how do they balance this interaction activity with the need to perform actions and look at objects relevant to their collaboration as well as multi-task on other duties and pay attention to the changing physical world around them? Claims that embodied 2D or 3D agents need not perform exactly as people are vacuous, because the engagement process is critical for people to know that they are being understood (the backward-looking aspect of engagement) and whether they intend to continue their interaction (the forward-looking aspect).

The approach at MERL to answering the above questions involves a three-pronged attack: study of human-human data to understand what people do in conversational collaboration, design of a robot that can produce such behaviour and interpret behaviours from a human collaborator, and evaluation of the robot's behavior with human users. What is learned from these efforts then allows the group to (1) refine the models of interaction that the robot uses and (2) include new interpretation and generation capabilities in the Collagen middleware system. Recent papers (listed below) provide details of what has been learned from human-human data, creating a robot architecture, and results of studies of people interacting with the robot. More about Collagen and human robot interaction can be found on the web sites below.

FOR INFORMATION

Candace Sidner is Senior Research Scientist at teh Mitsubishi Electric Research Labs.
Email: sidner@merl.com
Web: www.merl.com/people/sidner/
Project web sites:
www.merl.com/projects/FormsTalk/
www.merl.com/projects/collagen/
www.merl.com/projects/hosting/

References

[1] DeKoven, E., Keyson, D. V., Freudenthal, A., 'Designing Collaboration in Consumer Products' (Series 2)(abstract). In: *ACM conference on Computer Human Interaction (CHI 2001)*, pp. 127-132. ACM Press. Seattle, Washington, 2001.

[2] Rich, C., Sidner, C.L., 'COLLAGEN: A Collaboration Manager for Software Interface Agents', *An International Journal: User Modeling and User-Adapted Interaction*, Vol. 8, Issue 3/4, pp. 315-350, 1998.

[3] Rickel, J., Lesh, N. B., Rich, C., Sidner, C.L., Gertner, A., 'Collaborative Discourse Theory as a Foundation for Tutorial Dialogue', International Conference on Intelligent Tutoring Systems (ITS), *Lecture Notes in Computer Science*, Vol 2363, pp. 542-551, Springer Verlag, 2002.

[4] Sidner, C.L., Forlines, C., 'Subset Languages for Conversing with Collaborative Interface Agents', *International Conference on Spoken Language Processing (ICSLP)*, September 2002.

[5] Sidner, C.L., Lee, C., Lesh, N., 'Engagement by Looking: Behaviors for Robots when Collaborating with People', *DiaBruck: Workshop on the Semantics and Pragmatics of Dialogue (DiaBruck 2003)*, September 2003.

[6] Sidner, C.L., Kidd, C.D., Lee, C., Lesh, N. 'Where to look: a study of human-robot interaction', in *Proceedings of the Intemational Conference in Intelligent User Interfaces*, ACM Press, 2004, forthcoming.

MEANING: Developing Multilingual Web-scale Language Technologies

German Rigau, *Universitat Politècnica de Catalunya*

Progress is being made in Natural Language Processing (NLP) but there is still a long way towards Natural Language Understanding (NLU). An important step towards this goal is the development of technologies and resources that deal with concepts rather than words.

Even now, building large and rich knowledge bases takes a great deal of expensive manual effort; this has severely hampered Human Language Technologies (HLT) application development. For example, dozens of person-years have been invested into the development of wordnets for various languages, but the data in these resources are still not sufficiently rich to support advanced concept-based HLT applications directly. Furthermore, resources produced by introspection usually fail to register what really occurs in texts. Applications will not scale up to working in the open domain without more detailed and rich general-purpose and also domain-specific linguistic knowledge. However, progress is difficult because to be able to build the next generation of intelligent open domain HLT application systems we need to solve two complementary intermediate tasks: Word Sense Disambiguation (WSD) and large-scale enrichment of Lexical Knowledge Bases. Progress is thus hampered by the following paradox:

1) In order to achieve accurate WSD, we need far more linguistic and semantic knowledge than is available in current lexical knowledge bases (e.g, current wordnets).

2) In order to automatically enrich Lexical Knowledge Bases we need to acquire information from corpora which have been accurately tagged with word senses.

Providing innovative technology to solve this problem is one of the main challenges of the MEANING project (Rigau et al. 2002).

In order to extend the state-of-the-art in HLT, MEANING is promoting research on (1) innovative processes and tools for automatic acquisition of lexical knowledge from large-scale document collections; (2) novel techniques for accurately selecting the sense of open-class words in a large number of languages; (3) ways to enrich existing multilingual linguistic knowledge resources by automatically mapping information across languages.

German Rigau

1) Dealing with knowledge acquisition

The acquisition of linguistic knowledge from corpora has been a very successful line of research. However, much of the use of the acquired knowledge has been hampered by the fact that the texts are not sense-disambiguated, and therefore, only certain knowledge about words can be acquired; that is, subcategorisation, selectional preferences, etc. It is a well-established fact that much of the linguistic behaviour of words can be better explained with reference to word senses.

Having texts automatically sense-tagged with high accuracy will produce significantly better knowledge, including subcategorisation frequencies, domain information, topic signatures, selectional preferences, specific lexico-semantic relations, thematic role assignments, and diathesis alternations. Furthermore, it will allow the investigation of automatic methods for dealing with new senses not present in current wordnets and clustering of word senses.

2) Dealing with concepts

Word sense disambiguation (WSD) is the task of assigning the appropriate meaning (sense) to a given word in a text or discourse. Word sense ambiguity is a central problem for many established HLT applications (for example machine translation, information extraction, and information retrieval). This is also the case for associated sub-tasks (e.g., reference resolution and parsing). For this reason many international research groups

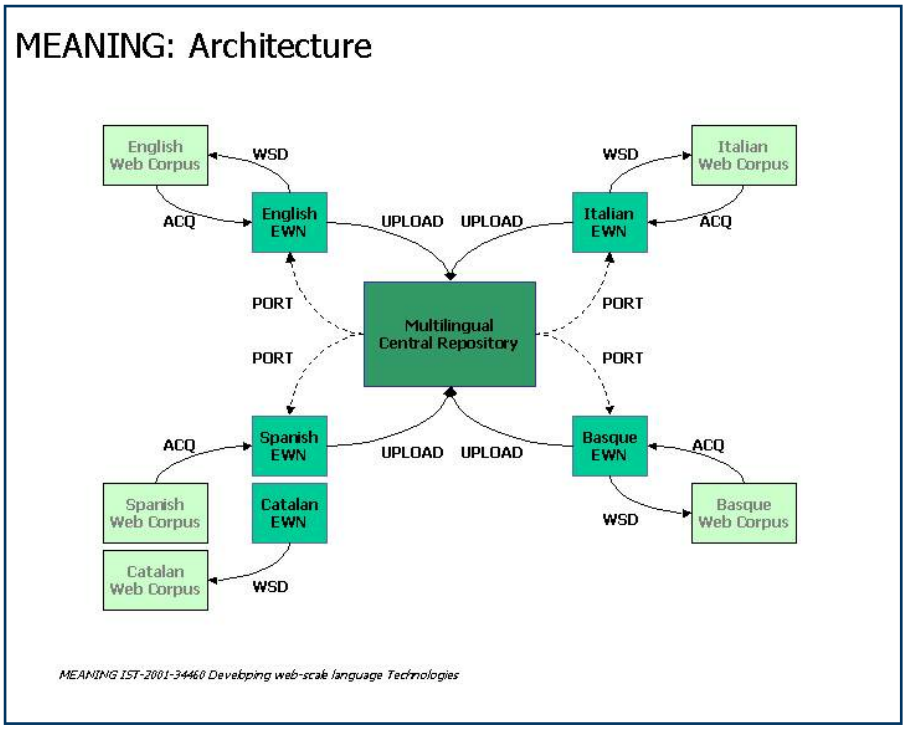

are working on WSD, using a wide range of approaches. However, no large-scale broad-coverage accurate WSD system has been built to date. With current state-of-the-art accuracy in the range 60-70% (Florian et al. 2002), WSD is perhaps one of the most important open problems in NLP.

A promising current line of research uses semantically annotated corpora to train machine learning (ML) algorithms to decide which word sense to choose in which contexts. The words in these annotated corpora are tagged manually with semantic classes taken from a particular lexical semantic resource (most commonly WordNet). Many standard ML techniques have been tried, such as Bayesian learning, exemplar based learning, decision lists, and recently margin-based classifiers like boosting and support vector machines. These approaches are termed "supervised" because they learn from previously sense annotated data and therefore they require a large amount of human intervention to annotate the training data.

However, supervised WSD systems suffer from the "knowledge acquisition bottleneck": it takes them mere seconds to process the training materials that take months to annotate manually. So, although machine learning classifiers are undeniably effective, they are not feasible until we can obtain reliable unsupervised training data. Ng (1997) estimates that the manual annotation effort necessary to build a broad-coverage word-sense-annotated English corpus is about 16 person-years; and this effort would have to be replicated for each different language. Unfortunately, many people think that Ng's estimate might fall short, as the annotated corpus thus produced is not guaranteed to enable high accuracy WSD.

Some recent work focuses on reducing the acquisition cost and the need for supervision in corpus-based methods for WSD. Leacock et al. (1998) and Mihalcea and Moldovan (1999) automatically generate arbitrarily large corpora for unsupervised WSD training, using the synonyms or definitions of word senses provided in WordNet to formulate search engine queries over the Web. In another independent research area Yarowsky, (1995) and Blum and Mitchell, (1998) have shown that it is possible to reduce the need for supervision with the help of large amounts of unannotated data. Following these ideas, MEANING is developing knowledge-based prototypes for obtaining accurate examples from the web for specific WordNet synsets, as well as large quantities of unannotated examples.

But in order to make significant advances in WSD system accuracy, systems need to be able to use types of lexical knowledge that are not currently available in wide-coverage lexical knowledge bases: for example subcategorisation frequencies for predicates (particularly verbs) rely on word senses or selectional preferences of predicates for classes of arguments, amongst others.

3) Dealing with multilingualism

In MEANING, the idiosyncratic way the meaning is realised in a particular language is captured and ported to the rest of languages involved in the project using the EuroWordNet architecture (Vossen 1998). In EuroWordNet local wordnets are linked via an Inter-Lingual-Index (ILI) allowing the connection from words in one language to translation-equivalent words in any of the other languages. In that way, technological advances in one language can help the other languages connected. For instance, for Basque, being an agglutinative language with very rich morphological-syntactic information, it is possible to extract semantic relations that would be more difficult to capture in other languages. However, Basque is not as widely found on the web as the others. Using this approach it is possible to balance both gaps.

Although the technology to provide compatibility across wordnets exits (Daude et al., 2003), new research is being carried out in MEANING for porting and uploading the various types of knowledge across languages, and new ways to validate the ported knowledge in the target languages.

The three research lines mentioned above have been explored separately with relative success. In fact, until now,

no research group in isolation has tried to combine all these factors. We designed the MEANING project convinced that only a combination of all relevant knowledge and resources will be able to produce significant advances in this crucial research area.

MEANING is now performing an innovative bootstrapping process (see figure 1) to deal with the inter-dependency between WSD and knowledge acquisition. This bootstrapping process consist of three consecutives cycles, each one including large-scale WSD, acquisition, and porting. In each cycle, highly accurate WSD systems allow more accurate acquisition of knowledge, which is placed into a common Multilingual Central Repository based on EuroWordNet (that will help future WSD and acquistion phases).

A wide range of techniques are required to progressively automate the knowledge lifecycle – in particular, for developing a trustworthy semantic web infrastructure and a multilingual ontology framework to support advanced knowledge management. There is no doubt that these processes will require extracting high-level meaning from the large collections of content data and their automatic representation in a common knowledge base. This is the main research goal of the MEANING project.

FOR INFORMATION

German Rigau is Associate Professor at the Universitat Politècnica de Catalunya.
Email: rigau@lsi.upc.es, **Web:** www.lsi.upc.es/~rigau/

The partners in the MEANING project are: Universitat Politecnica de Catalunya, ITC-IRST, University of the Basque Country, University of Sussex, Reuters Limited, and Irion Technologies B.V.

Project web site: www.lsi.upc.es/~nlp/meaning/meaning.html

References

Blum A. and Mitchel T. 'Combining labelled and unlabeled data with co-training'. In *Proceedings of the 11th Annual Conference on Computational Learning Theory*. 1998.

Daude J., Padro L. and Rigau G. 'Validation and Tuning of Wordnet Mapping Techniques'. In *Proceedings of the International Conference on Recent Advances in Natural Language Processing (RANLP'03)* Borovets, Bulgaria, 2003.

Florian R., Cucerzan S., Schafer C. and Yarowsky D. 'Combining Classifiers for Word Sense Disambiguation'. *Natural Language Engineering* 1(1) pp. 1-14. 2002.

Leacock, C. Chodorow, M. and Miller, G.A. 'Using Corpus Statistics and WordNet Relations for Sense Identication', *Computational Linguistics*, 24(1), 1998.

Mihalcea R. and Moldovan D. 'An automatic method for generating sense tagged corpora'. In *Proceedings of American Association for Artificial Intelligence*. 1999.

Ng. H. T. 'Getting Serious about Word Sense Disambiguation'. In *Proceedings of Workshop "Tagging Text with Lexical Semantics: Why, what and how?"*, Washington, 1997.

Rigau G., Magnini B., Agirre E., Vossen P. and Carroll J., 'MEANING: A Roadmap to Knowledge Technologies'. *Proceedings of COLING Workshop "A Roadmap for Computational Linguistics"*. Taipei, Taiwan. 2002.

Vossen P. *EuroWordNet: A Multilingual Database with Lexical Semantic Networks*, Kluwer Academic Publishers, Dordrecht. 1998.

Yarowsky D., 'Unsupervised word sense disambiguation rivaling supervised methods'. In *Proceedings of the 33rd Annual Meeting of the Association for Computational Linguistics*. 1995.

'Placa de Catalunya', 'La Rambla', 'The Olympic Quarter', and the pilgrims' place 'Montserrat', a bit out of town.

In short, the students enjoyed their time in beautiful Barcelona, the city, the courses, the intercourse with their fellow students and lecturers. It is commonly agreed that it was a good and fruitful experience. Besides, all participants share the hope for a continuation of the summer school tradition, thus providing other students with similar opportunities in the future, and would like to thank the organisers for all their care and help before and during the organisation of the school.

FOR INFORMATION

Petra Wagner is a lecturer at the Institut für Kommunikationsforschung und Phonetik at the Universität Bonn

Email: wagner@ikp.uni-bonn.de

Web: www.ikp.uni-bonn.de/~pwa/

cont'd from p. 10

Workshop report

The EMLS Summer School 2003 in Barcelona

compiled by ***Petra Wagner*** *based on reports by* ***Jayme Blans, Gina Koobs,*** *and* ***Khiet Truong*** *(all Utrecht), and* ***Corinna Kann*** *and* ***Eva Lasarcyk*** *(both Bonn)*

This year's Summer School of the European Masters in Language and Speech took place from July 7-11 at UPC Barcelona. It was the fourth of its kind, and once again it was a success with both students and lecturers. The summer school has become increasingly popular in recent years. In July 2003, it attracted more than 60 participants from all over Europe to the beautiful Catalan capital for a week dedicated to speech processing and computational linguistics. A wide range of tutorials covering a lot of different topics, student presentations, and posters were awaiting the participants, and provided a full one week programme.

The following text presents the impressions of several students from the Netherlands and Germany who hopefully are representative for the majority of participants.

The academic programme of the school included twelve tutorials, ranging from speech processing issues like speech synthesis and dialogue systems, to "knowledge representation" issues like natural language processing with PROLOG or building semantic networks. Most of the student presentations concentrated on ongoing or recently finished M.A. projects. The students presented their work either as a 30-minute talk or with the help of prepared posters.

The students especially enjoyed the hands-on approaches in many courses such as Building Corpora and Limited Domain Synthesis Exercise, since these helped to get a better understanding of all the necessary steps involved building working applications or databases. For prospective students of upcoming summer schools, it should be noted that the courses are introductions to the work of the lecturers rather than a substitute for a one-semester course held at university. Given an average of seven hours duration, a summer school course cannot provide one with a detailed introduction as well as developing time for a full project. Consequently, the participants were not expected to develop complex code and algorithms by themselves, but rather copy-paste and modify existing ones. In general, less time was spent on explaining details of the programming code but rather on the general procedure of the individual approach or application. The courses thus provide the students with a basis on which they develop their own individual applications.

While it was sometimes hard to gather information about course contents in advance of the summer school, most questions were answered via email immediately. Once people had arrived in Barcelona, the organisational aspects and interchange of information went smoothly. The accommodation was pleasant for several reasons: it was located in a nicely equipped Youth Hostel in the centre of Barcelona. Staying in the youth hostel gave every participant a chance to meet some of the other European students and learn more about their ongoing projects; and besides, it was close to the beach which turned out to be the students' favourite spot! At anytime of the day (or night) it proved to be an excellent point to meet other EMLS students in an informal atmosphere and provided the best environment to relax from the lectures. Unlike that of the students, the lecturers' accommodation was close to the university campus outside the city centre. The students on the whole appreciated the fact that it was not the other way round.

Besides doing a good job of preparing the academic programme of the summer school, the local organisers also surprised the participants with a pleasant excursion to Parc Güell, which was generally enjoyed. One got the chance to socialise with other participants and lecturers, of exchanging ideas, and simply absorbing the historical and architectural aspects of the park. Furthermore, the welcoming reception at the impressive Barcelona Town Hall was an amazing experience, making every participant feel honoured to have been treated in such a spectacular and friendly fashion.

Despite the day-filling EMLS programme, enough time was left for individual exploration of the town with its interesting and famous attractions and its very own 'flair'. The most highly frequented place was probably the church 'Sagrada Familia' by Gaudì, which is still in the process of being built. Other places of interest were

cont'd on p. 9

Participants at the EMLS Summer School

Natural Language Components for Smart Applications

***Atro Voutilainen**, Connexor*

Informed decision-making in corporations and other environments needs access to relevant knowledge, primarily coded as natural language text. Potentially relevant text is accessible in large and growing volume over the Web. Text volume and variety is too large for casual reading, so help is needed to find the relevant information and serve it to the consumer in an appropriate form and language. A software industry, Knowledge Management, has emerged to meet this knowledge need.

Because knowledge is largely coded as free-form text, a KM application should be able to find and use the information thus 'hidden' in text, i.e, analyse, and to some extent understand, natural language. Despite the half-century research tradition in natural language processing, automating language analysis and understanding are still regarded as very hard challenges. There is a demand for NLP expertise in the KM industry, though maybe not yet as high as one versed in language technology might consider reasonable.

Connexor was started (originally under the name `Conexor') in 1997 by three former members of the University of Helsinki's NLP group – Pasi Tapanainen, Timo Järvinen and Atro Voutilainen – after nearly two decades of successful academic research into linguistically oriented natural language processing (including finite state morphology, tagging based on linguistic constraints, functional dependency syntax) at the University of Helsinki and other research centres. During its first years Connexor mainly developed products for the end-user market (e.g., a grammar/style checker for English known as TrueStyler; a terminology management and content navigation program called NaviTerm). After the turn of the century, Connexor chose to focus on developing and licensing its language analysis technology as embeddable components to application developers. Connexor's 'Machinese' products enrich natural language texts (currently with support for ten languages: English, French, Spanish, German, Dutch, Italian, Finnish, Swedish, Norwegian, Danish) with an interpretative layer served as a programmer's interface (API) to application developers working with various platforms (e.g., Windows, Linux, Unix) and technologies. Using the Machinese API, application developers can access and use relevant information 'decoded' by Machinese, without getting too involved in the complexities of natural language and linguistic analysis.

Connexor's language processing technology is currently used in some 150 software houses and R&D labs worldwide, both academic (e.g., Massachusetts Institute of Technology, Harvard University, New York University, University of Tokyo) and commercial, mainly for the following application areas (some commercial organisations that we are permitted to mention by name are shown in parentheses):

- knowledge management (ClearForest, Fast Search and Transfer, Basis Technology, General Electric R&D, Corpora plc, Microsoft);
- machine translation (Mimos Berhad, Kielikone, Langsoft, Master's Innovations);
- education (Bookette, VanLindon Training Methods, Kone Elevators);
- speech (Toshiba Corporation).

Design of language models is knowledge-based and linguistically oriented rather than stochastic. Linguistic insights are refined and empirically validated against extensive corpus data. The linguistic approach developed by Connexor's founders and their colleagues has shown its success in accurate and informative language analysis, as can be seen in their many publications in leading language engineering forums (e.g., conference series organised by ACL and ICCL).

Let us now look more closely at the Machinese products. The architecture and compilers of Connexor's products are based on Connexor's finite state calculus and language-generic design. Adding support to a new language basically means building appropriate language models (extensive morphosyntactic lexicons and grammars) for the language.

1	I	I	subj>2	@SUBJ %NH PRON PERS NOM SG1
2	remember	remember	main>0	@+FMAINV %VA V PRES
3	having	have	v-ch>4	@-FAUXV %AUX ING
4	seen	see	obj>2	@-FMAINV %VA EN
5	her	she	obj>4	@OBJ %NH PRON PERS ACC SG3
6	somewhere	somewhere	loc>4	@ADVL %EH ADV
7	before	before	tmp>4	@ADVL %EH ADV
8	.	.		

Figure 1: Machinese Syntax analysis of "I remember having seen her somewhere before."

Machinese Phrase Tagger

Machinese Phrase Tagger is a fast light parser that enriches text with base forms and tags for part of speech, morphology, and basic meaningful entities (e.g., modifier-head sequences), i.e., it produces word class and shallow syntactic tags. The sentence "They look at the different labour market models of Europe and the US" gives the following output when analyzed with Machinese Phrase Tagger:

They	they	@NH PRON
look	look	@MAIN V IND PRES
at	at	@PREMARK PREP
the	the	@PREMOD DET
different	different	@PREMOD A
labour	labour	@PREMOD N
market	market	@PREMOD N
models	model	@NH N PL
of	of	@POSTMOD PREP
Europe	Europe	@NH N Prop
and	and	@CC CC
the	the	@PREMOD DET
US	US	@NH N Prop
.	.	

On the basis of this kind of simple markup, several useful tasks can be accomplished. For instance, terms and other content-rich expressions can be extracted by matching words and word sequences with appropriate word class and syntactic tags (e.g. adjectives or nouns as premodifier, followed by noun as nominal head). Basic phrases can also help to predict intonation, so Machinese Phrase Taggers are used in text-to-speech synthesis products.

Machinese Syntax

Machinese Syntax is a full syntactic parser that produces both morphological information for word-form tokens and functional dependencies representing relational information in sentences. For instance, Machinese Syntax shows simple and complex entities in sentences, and it describes relations between these entities:

- objects and ontological facts (names, organisations and places)
- actions ('who did what to whom'), and
- circumstances (where, when, why, how ...)

This makes Machinese Syntax a natural choice for knowledge-intensive applications.

Machinese Syntax output contains five fields:

- word position
- word-form
- base-form (a.k.a. lemma)
- functional dependency
- functional tag, surface-syntactic tag, and morphological tags

Consider the analysis of the sentence "I remember having seen her somewhere before" below. The fifth word in the sentence, "her", has the base form "she". It is analysed as a direct object (obj), dependent of word number 4. The functional dependency type is shown by the functional tag @OBJ, followed by the surface-syntactic label %NH indicating that the word is the head of the phrase. Morphological labels further specify that the word is the singular third-person form of a personal pronoun in the accusative case.

The explicit relational information is represented in the fourth field by functional dependencies. Functional dependencies consist of the function type label and a numerical index pointing to its head. In the sentence above, the main element of the sentence is the word numbered 2, which is expressed by the value 0 of the numerical index. The main element, "remember", has two direct dependents, the words "I" and "seen". The word "seen" is labelled as an object because it is the head of a participial clause acting as object of "remember", a complex form containing an auxiliary predicate "having" that is analysed as a verb chain element (v-ch) dependent on the main predicate. In addition, the verb "seen" has three direct dependents: the object "her", the adverbial "somewhere" expressing location, and the adverbial "before" expressing temporal relation. In this case, the semantic distinction between a locative and a temporal adverbial is expressed at the level of functional dependencies.

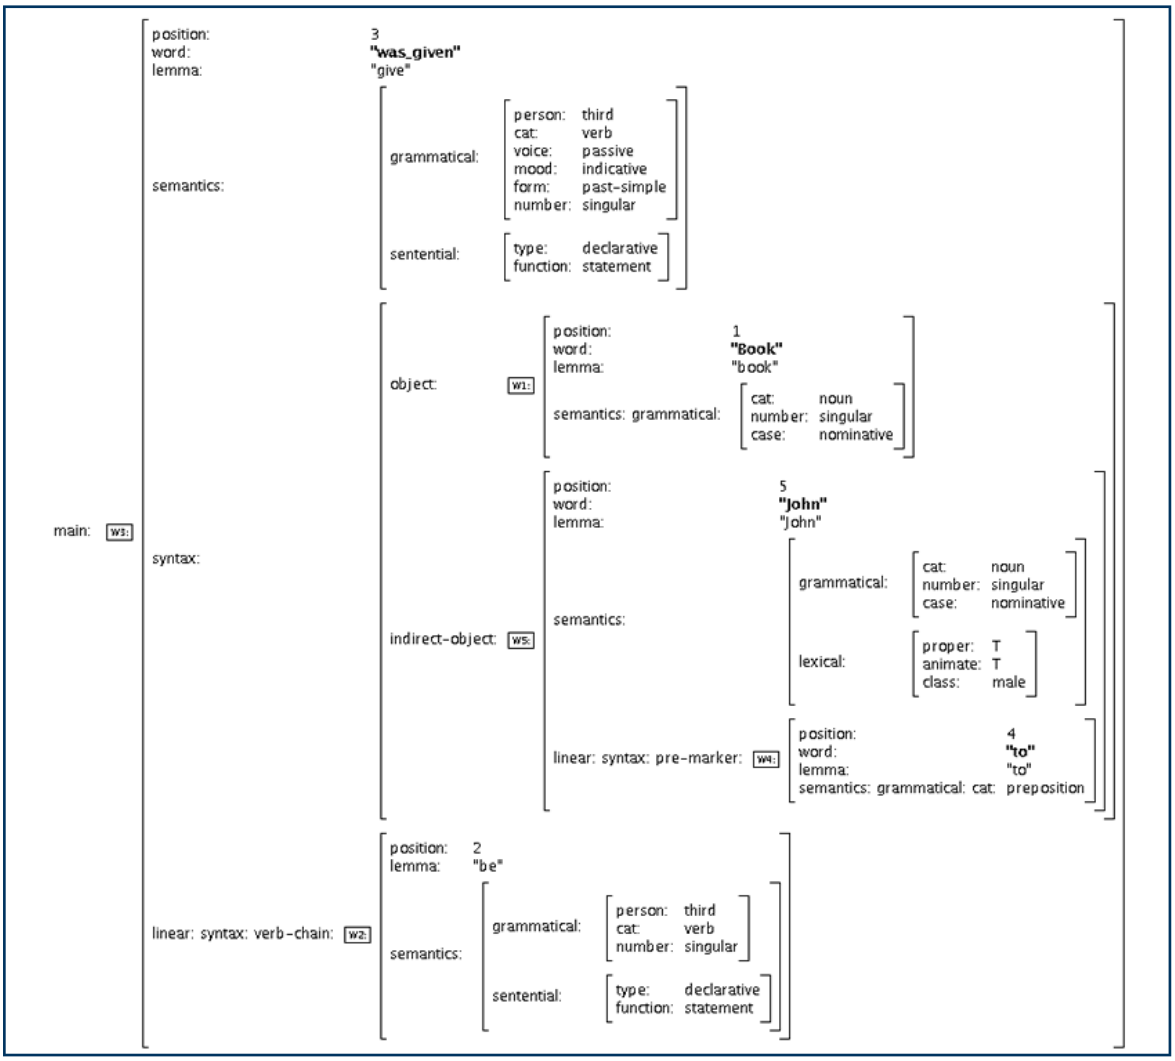

Figure 2: Machinese Semantics analysis of "A book was given to John"

Machinese Semantics

Machinese Semantics is a semantic analyser that provides semantic role recognition as well as grammatical, lexical, and sentential semantic features. These include, for example:

- recognition of multi-word forms as one entity (e.g., would_have_been_informed)
- harmonisation of the syntactic structures
- name recognition and classification (person, location, organisation)
- semantic feature recognition (human, animate, tool, durative, etc.)

It includes features that make it especially suitable for use as a source language analyser in machine translation or in information extraction.

Machinese Semantics produces output as (possibly recursive) attribute-value pairs. Categories are values of the corresponding attribute. Each node in the analysis is a single word or a multi-word unit. The parser produces three attributes whose value is a string. 'Word' is the running text token. 'Lemma' is the base form of the nucleus. 'Head' is printed only in those multi-word units where it is different from the lemma. After that follow the syntax, semantics, and linear matrices.

The Machinese Semantics analysis of the sample sentence "A book was given to John." is shown in Figure 2 as a graphical feature structure presentation.

Machinese Metadata

Machinese Metadata is a tool for:

- Information retrieval: automatic indexing
- Summarization of the information content: automatic keyword and name recognition
- Classification of the documents: automatic keyword and name recognition
- Terminological work: automatic technical term extraction

Machinese Metadata describes content using single-word and multi-word keywords and names. It extracts keywords and names automatically from information sources and indicates what the sources are about. Keywords and names summarise the essential information in a document.

The recognition of keywords, proper names, and technical term candidates is based on automatic content analysis that includes part-of-speech tagging, lexical and syntactic analysis, and analysis of the position and distribution of words. The content analysis produces an estimation of essential information in a document and the term-likelihood of the term candidates.

Machinese Metadata technology is used for automatic news analysis, indexing, and summary generation.

Current developments

An ongoing effort at Connexor is the compilation of extensive name resources for use in contextual recognition and classification of names and other information-rich entities. A part of this effort already shows in the current version of the Machinese Phrase Tagger. A Machinese product for identifying and classifying names and other entities for all Machinese languages will be released during this year.

Atro Voutilainen

FOR INFORMATION

Atro Voutilainen is CEO of Connexor

Email: atro.voutilainen@connexor.com
Web: www.connexor.com

Invitation and Call for Papers

Third International Conference on Natural Language Generation

Careys Manor
Brockenhurst
New Forest
UK
14-16 July 2004

INLG is the leading international conference on research into natural language generation.

INLG04 will take place in the beautiful New Forest in the south of England. In keeping with the tradition established by past INLG conferences, Careys Manor is a secluded self-contained venue with outstanding food, rooms, conference and fitness facilities.

Submissions are invited for the main and student sessions (deadline Mar 12) and the poster session (deadline May 14) of the conference, on any aspect of natural language generation. Full details of the calls and submission details can be found on the conference website.

INLG04 is organised by ITRI, University of Brighton on behalf of SIGGEN, the special interest group on generation of the Association for Computational Linguistics (ACL).

inlg04@itri.brighton.ac.uk www.itri.brighton.ac.uk/inlg04

Calendar

Future Events

Feb 15-21 *Fourth International Conference on Intelligent Text Processing and Computational Linguistics (CICLing-2004):* Seoul, Korea
Email: gelbukh@CICLing.org URL: www/CICLingorg/2004

Mar 6-9 *Sixth Teaching and Language Corpora Conference (TaLC 2004):* Granada, Spain
Email: talc6@ugr.es URL: www.ugr.es/local/talc6

Mar 10-12 *Seventh International Conference on the Statistical Analysis of Textual Data:* Louvain-la-Neuve, Belgium
Email: fairon@tedm.ucl.ac.be URL: www.jadt.org

Mar 22-24 *First International Joint Conference on Natural Language Processing (IJCNLP-04):* Hainan Island, China
Email: ijcnlp04.enquiry@cityu.edu.hk URL: www.cipsc.org.cn/IJCNLP-04

Apr 13 *First International Workshop on Natural Language Understanding and Cognitive Science (NLUCS-2004):* Porto, Portugal
Email: workshops@iceis.org URL: www.iceis.org

Apr 19-22 *TALN04 (Traitement Automatique du Langage Naturel):* Fez, Morocco
Email: taln2004@lpl.univ-aix.fr URL: www.lpl.univ-aix.fr/jep-taln04

Apr 26-27 *European Association for Machine Translation Workshop:* Malta
Email: mike.rosner@um.edu.mt URL: www.eamt.org

Submission deadlines

Feb 22 *ESSLLI-2004 Stduent Session:* Nancy, France Aug 9-20,
Email: paulege@magic.fr, URL: esslli2004.loria.fr

Mar 8 *ACL04 Student research Workshop:* Barcelona, Spain, Jul 21-26,
Email: acl04-student@list.cs.brown.edu, URL:www.ad04.org

Mar 12 *Third International Conference on Natural Language Generation (INLG04):* New Forest, UK, Jul 14-16, Email: inlg04@itri.brighton.ac.uk, URL: www.itri.brighton.ac.uk/inlg04

Mar 12 *CATALOG'04 (Eighth Workshop on the Semantics and Pragmatics of Dialogue)*, Barcelona, Spain, Jun 19-21, Email: catalog04@upf.edu, URL:www.upf.edu/catalog

Mar 15 *Ninth International workshop on Speech and the Computer (SPECOM'2004):* St Petersburg, Russia, Sep 20-22, Email:specom@mail.iias.spb.su, URL: www.spiiras.nw.ru/speech

This is only a selection – see www.elsnet.org/cgi-bin/elsnet/events.pl for details of more events and deadlines.pl for more deadlines.

If you would like to write a review of any of these (or other language/speech related events you attend), please contact the *ELSNews* editor.

ELSNET

Office

Steven Krauwer,
Co-ordinator
Brigitte Burger,
Assistant Co-ordinator
Utrecht University (NL)

The ELSNET Participants: Academic Sites

A	University of Vienna
A	Austrian Research Institute for Artificial Intelligence (OFAI)
A	Vienna University of Technology
AU	Macquarie University
B	University of Antwerp - UIA
B	Katholieke Universiteit Leuven
BG	Bulg. Acad. Sci. - Institute of Mathematics and Informatics
BY	Belorussian Academy of Sciences
CH	SUPSI University of Applied Sciences
CH	University of Geneva
CZ	Charles University
D	Universitaet des Saarlandes
D	Ruhr-Universitaet Bochum
D	Universität des Saarlandes CS-AI
D	German Research Center for Artificial Intelligence (DFKI)
D	Institut für Angewandte Informationsforschung
D	Universität Erlangen-Nürnberg - FORWISS
D	Universität Hamburg
D	Christian-Albrechts University, Kiel
D	Universität Stuttgart-IMS
DK	University of Southern Denmark
DK	Center for Sprogteknologi
DK	Aalborg University
E	Universidad Politécnica de Valencia
E	University of Granada
E	Universidad Nacional de Educación a Distancia (UNED)
E	Polytechnic University of Catalonia
E	Universitat Autonoma de Barcelona
E	Universidad Politécnica de Madrid
EL	National Centre for Scientific Research (NCSR) 'Demokritos'
EL	University of Patras
EL	Institute for Language & Speech Processing (ILSP)
F	LORIA
F	Inst. National Polytechnique de Grenoble
F	LIMSI/CNRS
F	IRISA/ENSSAT
F	Université Paul Sabatier (Toulouse III)
F	Université de Provence
GE	Tbilisi State University, Centre on Language, Logic and Speech
HU	Lóránd Eötvös University
HU	Technical University of Budapest
I	Università degli Studi di Pisa
I	Consorzio Pisa Ricerche
I	Fondazione Ugo Bordoni
I	IRST
I	Consiglio Nazionale delle Ricerche
IRL	Trinity College, University of Dublin
IRL	University College Dublin
LT	Inst. of Mathematics & Informatics
NL	Foundation for Speech Technology
NL	University of Twente
NL	University of Groningen
NL	Tilburg University
NL	Eindhoven University of Technology (TUE)
NL	University of Nijmegen
NL	Leiden University
NL	Utrecht University
NL	Netherlands Organization for Applied Scientific Research TNO
NL	University of Amsterdam (UvA)
NO	Norwegian University of Science and Technology
NO	University of Bergen
P	University of Lisbon
P	INESC ID Lisboa
P	New University of Lisbon
PL	Polish Academy of Sciences
RO	Romanian Academy
RU	Russian Academy of Sciences, Moscow
S	KTH (Royal Institute of Technology)
S	Linköping University
TR	Sabanci University
UA	IRTC UNESCO/IIP
UK	University of Edinburgh
UK	Leeds University
UK	University of Sheffield
UK	University of Essex
UK	University College London
UK	The Queen's University of Belfast
UK	University of Brighton
UK	University of York
UK	UMIST
UK	University of Dundee
UK	University of Ulster
UK	University of Cambridge
UK	University of Sussex
UK	University of Sunderland

Industrial Sites

B	Dhaxley Translations
CH	Localization Industry Standards Association
D	Novotech GmbH
D	Sympalog Speech Technologies AG
D	DaimlerChrysler AG
D	Langenscheidt KG
D	Verlag Moritz Diesterweg GmbH
D	aspect Gesellschaft für Mensch-Maschine Kommunikation mbH
D	Philips Research Laboratories
D	Grundig Professional Electronics GmbH
D	Acolada Gmbh
D	IBM Deutschland
D	Varetis Communications
D	Heartsome Europe GmbH
D	Xtramind Technologies GmbH
D	Sympalog Voice Solutions GmbH
D	Scansoft Aachen GmbH
DK	Tele Danmark
DK	Zacco A/S
E	SchlumbergerSema sae
E	Telefonica I & D
EL	KNOWLEGDE S.A.
F	LINGA s.a.r.l.
F	Systran SA
F	Xerox Research Centre Europe
F	Memodata
F	Aerospatiale
F	VECSYS
F	SCIPER
F	TGID
FIN	Kielikone Oy
FIN	Nokia Research Center
HU	MorphoLogic Ltd.
I	OLIVETTI RICERCA SCpA
I	LOQUENDO
LV	TILDE
NL	Compuleer
NL	Knowledge Concepts BV
NL	Sopheon NV
NL	IP Globalnet Nederland BV
PL	Neurosoft Sp. z o.o.
RU	Russicon Company
RU	ANALIT Ltd
S	Sema Infodata
S	Telia Promotor AB
UK	Vocalis, Ltd.
UK	Hewlett-Packard Laboratories
UK	Canon Research Centre Europe Ltd
UK	ALPNET UK Limited
UK	Sharp Laboratories of Europe Ltd
UK	BT Adastral Park
UK	Logica Cambridge Ltd.
UK	20/20 Speech LTD
UK	Tangent Telecom Ltd
UK	Cambridge Algorithmica Limited
UK	Fourth Person Ltd.

What is ELSNET?

ELSNET is the European Network in Human Language Technologies. ELSNET is sponsored by the Human Language Technologies programme of the European Commission; its main objective is to foster the human language technologies on a broad front, creating a platform which bridges the gap between the natural language and speech communities, and the gap between academia and industry.

ELSNET operates in an international context across discipline boundaries, and deals with all aspects of human communication research which have a link with language and speech. Members include public and private research institutions and commercial companies involved in language and speech technology.

ELSNET aims to encourage and support fruitful collaboration between Europe's key players in research, development, integration, and deployment across the field of language and speech technology and neighbouring areas.

ELSNET seeks to develop an environment which allows optimal exploitation of the available human and intellectual resources in order to advance the field. To this end, the Network has established an infrastructure for the sharing of knowledge, resources, problems, and solutions across the language and speech communities, and serving both academia and industry. It has developed various structures (committees, special interest groups), events (summer schools, workshops), and services (website, e-mail lists, *ELSNews*, information dissemination, knowledge brokerage).

Electronic Mailing List

elsnet-list is ELSNET's electronic mailing list. Email sent to elsnet-list@let.uu.nl is received by all member site contact persons, as well as other interested parties. This mailing list may be used to announce activities, post job openings, or discuss issues which are relevant to ELSNET. To request additions/deletions/changes of address in the mailing list, please send mail to elsnet@let.uu.nl

Subscriptions

Subscriptions to *ELSNews* are currently free of charge. To subscribe, visit **http://www.elsnet.org** and follow the links to *ELSNews* and "subscription".

FOR INFORMATION

ELSNET
Utrecht Institute of Linguistics OTS, Utrecht University,
Trans 10, 3512 JK, Utrecht, The Netherlands
Tel: +31 30 253 6039
Fax: +31 30 253 6000
Email: elsnet@elsnet.org
Web: http://www.elsnet.org