

## First International Conference on Global WordNet

*Piek Vossen, Irion Technologies, Delft*

The first international conference on the Global WordNet was held in India earlier this year. The conference, which took place from 21-25 January, was organised by the Central Institute of Indian Languages (CIIL) in Mysore, together with the Global WordNet Association (GWA), the Indian Institute of Technology (IIT), Bombay, and the International Institute of Information Technology (IIIT), Hyderabad.

There have been various workshops on WordNet in the past, but this was the first time that a whole conference was dedicated to the

topic. Most of the previous workshops have taken place at ACL conferences in the USA and Europe, and have had a strong focus on the English WordNet and its usage in NLP. But this conference was different: it also assembled WordNet builders and WordNet tuners, especially for many other languages as well. The fantastic venue of India and the perfect professional organisation thus resulted in an excellent and exotic



*The perfect background for a discussion on cross-linguistic and cultural lexicalisation differences*

mixture for the 81 scholars from 19 countries who participated.

There was not just one theme or topic that dominated the discussion. Rather, the conference showed how widely WordNet is being used, both in regard to the range of types of applications (among which are language learning, text mining, named entity classification, and machine translation), and, even more clearly, in relation to the diversity of languages it now encompasses. WordNets are being developed and deployed in all kinds of languages: there

were representatives for Balkan languages, Russian, and further to the East, groups from Korea, China, and Japan, all working on WordNets in their own languages.

In a conference held in Mysore, the Indian languages were obviously well represented. In this respect, having the venue in India turned out to be very productive. India has the political will, well-trained linguists and

**GATE – New Release**  
Hamish Cunningham  
Kalina Bontcheva  
Diana Maynard  
Valentin Tablan 3

**MT Roadmap Workshop Report**  
Mikel Forcada  
Harold Somers 5

**Opinion Column**  
Kenneth Church 7

**SIG dial Page: Discourse and Dialogue Research in Japan**  
Yasuharu Den 8

**New ELSNET Members:**  
• **UNED**  
M.Felisa Verdejo 9  
• **Neurosoft**  
Cezary Grzegorz Dolega 10

**InSTIL 2002**  
Steve Larocca  
Kathleen Egan  
Fenfang Hwu  
Philippe Deldouque 11

**Speech Research and Technology in Kaunas University of Technology**  
Algimantas Rudzionis 13

**Future Events** 15

**Spring  
2002**

**elsnet**  
.....

engineers, and sufficient funding to carry out a long-term project such as the IndiaWordNet, in which WordNets for 20 Indian languages will be developed and interlinked. This conference brought together many Indian NLP people, linguists, and even some politicians who could be guided into making the right choices for supporting such a large project.



*Professor Udaya Narayana Singh, Director of CIIL, opening the Conference*

*To the left, B.D. Jayaram, the driving force behind the organisation of the conference; to the right, Piek Vossen and Christiane Fellbaum, on behalf of the Global Wordnet Association*

Spreading WordNets to other languages and cultures shows us, increasingly, the universal and non-universal aspects of WordNet. This will have important consequences for all cross-lingual information systems of the future. In this respect, it was very stimulating to see the work done on various Indian WordNets, among which even tribal languages were included. We even learned that clockwise and anti-clockwise direction of movement is an important lexicalised distinction in the tribal language Car Nicobarese.

Bringing together WordNet builders and language engineers indicated many new directions for research. We saw some excellent graphic representations and WordNet editors that will help the development and understanding of the complex structures. There were fundamental discussions on semantic relations (new and old) and WordNet's top levels. This also showed the strong need to further standardise the WordNet model.

The extension and tuning of WordNet to specific domains turned out to be a continuing theme that was already evident from previous WordNet workshops. Domain tuning goes hand-in-hand with corpus anchoring of WordNets as a resource. Domain corpora, and even the WWW as an overall corpus, form the basis of most of the current approaches for WordNet development. This seems to be the way to go in the future.

There were two tutorial sessions. One was given by Christiane Fellbaum, Piek Vossen, and Palmira Marrafa, and the other by Eneko Agirre and Geman Rigau. Both of these discussed various issues in relation to building WordNets for English and various other European languages.

The business meeting was open to all the conference participants and was conducted by three board members of the GWA, namely, Piek Vossen, Christiane Fellbaum, and Palmira Marrafa. The issues discussed during the business meeting were the activities and membership of the GWA, communication between members, standardisation of WordNets, and future meetings.

Finally, we would like to share with you a picture of the Global WordNet bus that took us to visit the most fantastic places and people around Mysore in South India



*The Global WordNet bus that drove us around Mysore*

We hope that the next conference, to be held in January 2004, will be as stimulating and exotic as the first. The Global WordNet Conference welcomes any bids from interested parties who would like to host it.

## FOR INFORMATION

**Piek Vossen** is currently Chief Technology Officer at Irion Technologies, Delft, in The Netherlands. He was formerly Project Manager of EuroWordNet, based at the University of Amsterdam

**Email:** [Piek.Vossen@irion.nl](mailto:Piek.Vossen@irion.nl)

More activities of the Global WordNet Association can be found at: [www.hum.uva.nl/~ewn/gwa.html](http://www.hum.uva.nl/~ewn/gwa.html)

ELSNNews is published at the School of Cognitive and Computing Sciences, University of Sussex. It is printed by the University of Sussex Print Unit.

Editor: Jenny Norris.

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

ISSN 1350-990X

© ELSNET 2002

FOR INFORMATION  
Contributions to  
ELSNNews, and corrections  
should be sent to:  
[jennyn@cogs.susx.ac.uk](mailto:jennyn@cogs.susx.ac.uk)

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

Material for the next  
issue is due:  
10 June 2002





# GATE – A New Release

*Hamish Cunningham, Kalina Bontcheva, Diana Maynard, & Valentin Tablan, University of Sheffield*

GATE – a General Architecture for Text Engineering – is an infrastructure for developing and deploying software components that process human language. Version 2 was released on March 14, 2002.



**GATE — General Architecture for Text Engineering**

GATE has been in development at the University of Sheffield since 1995 and has been used in a wide variety of research and development projects, including Information Extraction (IE) in English, Bulgarian, Romanian, Bengali, Greek, Spanish, Swedish, German, Italian, and French.

GATE is a tool for:

- scientists performing experiments that involve processing human language
- companies developing applications with language processing components
- teachers and students of courses about language and language computation.

The system aims to help these groups in three ways:

- by specifying an architecture, or organisational structure, for language processing software
- by providing a framework (or class library, or Software Development Kit – SDK), that implements the architecture and can be used to embed language processing capabilities in diverse applications
- by providing a development environment built on top of the framework, consisting of convenient graphical tools for developing many types of component.

The architecture exploits component-based software development, object orientation, and mobile code. The framework and development environment are written in Java and are available as open-source free software from the GATE web site [see *FOR INFORMATION*, page 4].

As an architecture, GATE suggests that the elements of software systems that process natural language can usefully be broken down into various types of component, known as *resources*. These components are reusable software chunks with well-defined interfaces; GATE resources are specialised types of Java Bean, and come in three flavours:

- **LanguageResources (LRs)** represent entities such as lexicons, corpora, or ontologies
- **ProcessingResources (PRs)** represent entities that are primarily algorithmic, such as parsers, generators, or n-gram modellers

- **VisualResources (VRs)** represent visualisation and editing components that participate in GUIs.

(In practice, these definitions may be blurred if necessary for the particular application.)

Collectively, the set of resources integrated with GATE is known as a CREOLE: a Collection of REusable Objects for Language Engineering. All the resources are packaged as Java Archive (or 'JAR') files, with some additional XML configuration data. The JAR and XML files are made available to GATE by putting them on a web server, or simply by placing them in the local file space.

When using GATE to develop language processing functionality for an application, the developer uses the development environment and the framework to construct resources of the three types. This may involve programming, or the development of LRs, such as grammars, that are used by existing PRs, or a mixture of the two. The development environment is used for visualisation of the data structures produced and consumed during processing, and for debugging, performance measurement, and so on. When an appropriate set of resources has been developed, these resources can then be embedded in the target client application using the GATE framework (which is supplied as two JAR files).

GATE includes resources for common language engineering data structures and algorithms, including documents, corpora, various annotation types, a set of language analysis components for IE, and a range of data visualisation and editing components. GATE supports documents in a variety of formats, including XML, RTF, email, HTML, SGML, and plain text. In all cases the format is analysed and converted into a single unified model of annotation, compatible with other formats such as XCES, ATLAS, etc.

A family of PRs for language analysis is included in the shape of ANNIE, A Nearly-New Information Extraction system. These components use finite-state techniques to implement various tasks – from tokenisation to semantic tagging or verb phrase chunking. ANNIE is currently in active use in a number of industrial and academic projects.

Feature

**Spring  
2002**

**elsnet**  
.....

There are three other facilities in GATE which deserve special mention:

- **JAPE** – a Java Annotation Patterns Engine – provides pattern/action rules over annotations based on regular expressions
- The “**annotation diff**” tool in the development environment – this implements performance metrics such as precision and recall for comparing annotations and for quantitative evaluation purposes
- **GUK** – the GATE Unicode Kit – fills in some of the gaps in Java’s support for Unicode, e.g., by adding input methods for various languages, from Urdu to Chinese.

### An Example of GATE in Use

Let’s imagine that a developer called Fatima is building an email client for Cyberdyne Systems’ large corporate Intranet. In this application she would like to have a language processing system that automatically spots the names of people in the corporation and transforms them into *mailto* hyperlinks.

A little investigation shows that GATE’s existing components can be tailored to this purpose. Fatima starts up the development environment, and creates a new document containing some example emails. She then loads some processing resources that will do named-entity recognition (a tokeniser, gazetteer, and semantic tagger), and creates an application to run these components on the document in sequence. Having processed the emails, she can see the results in one of several viewers for annotations.

The GATE components are a decent start, but they need to be altered to deal specifically with people from Cyberdyne’s personnel database. Therefore, Fatima creates new “cyber-” versions of the gazetteer and semantic tagger resources, using the “bootstrap” tool. This tool creates a directory structure on disk that has some Java stub code, a Makefile, and an XML configuration file. After several hours struggling with badly written documentation, Fatima manages to compile the stubs and create a JAR file containing the new resources. She tells GATE the URL of these files and the system then allows her to load them in the same way that she loaded the built-in resources earlier on.

Fatima then creates a second copy of the email document, and uses the annotation-editing facilities to mark up the results that she would like to see her system producing. She saves this and the version on which she ran GATE, into her Oracle datastore. From now on she can follow this routine:

- Run her application on the email test corpus
- Check the performance of the system by running the “annotation diff” tool to compare her manual results with the system’s results. This gives her both percentage accuracy figures and a graphical display of the differences between the machine and human outputs
- Make edits to the code, pattern grammars, or gazetteer lists in her resources, and re-compile where necessary
- Tell GATE to re-initialise the resources
- Go back to the beginning.

To make the alterations that she requires, Fatima re-implements the ANNIE gazetteer so that it regenerates itself from the local personnel data. She then alters the pattern grammar in the semantic tagger to prioritise recognition of names from that source.

Eventually the system is running nicely. Now Fatima stops using the GATE development environment and works instead on embedding the new components in her email application. This application is written in Java, so embedding is very easy – the code for talking to GATE takes up only around 150 lines of the eventual application, most of which is just copied from the example in the StandAloneAnnie class on the web site.

Because Fatima is worried about Cyberdyne’s unethical policy of developing Skynet to help the large corporates of the West strengthen their stranglehold over the world, she wants to get a job as an academic instead (so that her conscience will only have to cope with the torture of students, as opposed to humanity). She takes the accuracy measures that she has attained for her system and writes a paper for the *Journal of Nasturtium Lqarithm Enâtement* describing the approach used and the results obtained. Because she used GATE for development, she can cite the repeatability of her experiments and offer access to example binary versions of her software by putting them on an external web server.

### FOR INFORMATION

The authors are all based at the Department of Computer Science of the University of Sheffield, UK. Hamish is Senior Research Scientist, and Kalina, Diana, and Valentin are all Research Associates

**Email:** hamish, kalina, diana, valyt@dcs.shef.ac.uk

**Web:** www.dcsshef.ac.uk/~hamish, kalina, diana, valyt

**GATE web site:** <http://gate.ac.uk/>

e

# Second ELSNET MT Roadmap at TMI in Japan

Workshop Report

*Mikel Forcada, Universitat d'Alacant, Spain, & Harold Somers, UMIST, Manchester, England*

As part of the 9th TMI Conference held at Keihanna, Kyoto, ELSNET Coordinator Steven Krauer organised a follow-up to the successful first MT Roadmap Workshop held at MT Summit VIII last September. Attended by about 25 participants, the one-day workshop consisted of invited lectures, refereed papers, and brainstorming.

In his opening address Krauer described ELSNET, particularly for the benefit of Japanese participants, and explained more precisely the purpose of the Roadmap workshops. He emphasised the interest in collecting Asian views on the future of MT.

The programme began with a thought-provoking invited talk from Professor Satoru Ikehara of Tottori University: *Toward the Realisation of Typological Semantic Pattern Dictionaries for MT*. His talk focussed on the need for "semantic pattern dictionaries" to improve lexical disambiguation. Ikehara criticised approaches to MT that relied only on syntax, suggesting that a dictionary of at least 100,000 semantic patterns, such as those found in his *Nihongo Goi-Taikei* lexicon (1997), could capture distinctions of meaning needed for good quality translation into English.

The next session saw three refereed papers. In the first, TMI host Francis Bond of NTT gave an interesting talk entitled *Toward a Science of Machine Translation*. Bond invited us to look to (human) translation theory, in the form of Nida's (1966) description of the human translation process. Noting that only one of Nida's nine steps in the translation process actually involves translation *per se*, Bond suggested MT counterparts for some of the other steps. For example, MT could involve multiple passes over the text, to identify the domain, isolate named entities, and locate passages where comparison with existing translations might be especially fruitful. Bond also suggested verifying the readability of the target text, and the addition of parenthetical explanations for unfamiliar local terms – a feature absent from current MT systems.

Mikel Forcada's talk focussed on the use of a simple finite-state direct translation approach based on aligned bi-texts automatically extracted from the Web. He suggested this as a way of rapidly developing MT systems for new language pairs, based on the principle of translation memory, but extending below the sentence level.

The final talk of the morning session was given by Nigel Ward of Tokyo University, with the intriguing title *Machine Translation in the Mobile and Wearable Age*. Ward described preliminary experiments with an apparatus involving a speech synthesiser, attached to a bilingual phrase-book, driven by a menu projected onto glasses worn by the user (see picture). A more appropriate title would perhaps have been "No Machine Translation ...", given his view that MT quality was too poor for this task and that it could be avoided by focussing on the human interaction aspects of communication. Not surprisingly, Ward's presentation provoked heated discussion.



*Jani Patokallio demonstrates a prototype wearable talking phrase book, developed at Tokyo University*

The afternoon session was devoted to brainstorming. Steven Krauer described current trends towards large MT projects which have the advantage of sharing goals, tools, resources, and technologies, but seem less flexible, with little emphasis on basic research.

Integrating MT into other applications is also favoured. Krauer suggested that MT is moving away from "one solution for all situations" towards "for each situation a solution". Interestingly, despite the opposite trend in NLP, the majority (by more than 3:1) of papers at recent MT conferences described rule-based rather than empirical systems, and he expressed his regret at the absence of breakthroughs in the last five years.

Spring  
2002

elsnet  
.....

>



The remainder of the workshop – including a humorous report from TMI 2012 (*sic*) by Harold Somers as light relief – involved general discussion in which participants were invited to specify a €100m single-goal five-year research project. Three perspectives were specified: basic research, industrial R&D, and users. The range of suggestions was very broad – some conservative, some far-fetched.

The following basic research topics were suggested:

- Formal theory of (machine) translation
- Semantic theory and/or universal meaning representation
- Theory of (cross-lingual) communication
- Eliminating the linguistic knowledge acquisition bottleneck
- Making large-scale language resources shareable (an issue of intellectual and industrial property rights)
- Integrating translation memory and MT in a single product
- Cross-training linguists and computer scientists
- A huge, massively annotated multitext aligned at various levels
- Robust speech processing including varying accents, language switching, and non-linguistic aspects such as facial and hand gestures.

R&D and user-motivated topics included:

- Plug-ins and services for mobile phones, for a variety of communication functions
- MT systems acquired from language learning texts
- Support for “technologically challenged” languages
- Focus on the “hard” part of cross-lingual communication (e.g., cultural differences)

- Evaluation and learning from user feedback (e.g., post-editing tools)
- User models
- Translating glasses to translate any text seen through them
- Multilingual web search and translation
- Translating stenography for live TV, conferences, etc.

An official summary of the workshop can be found at [www.elsnet.org/roadmap-tmi2002.html](http://www.elsnet.org/roadmap-tmi2002.html). Participants will no doubt look forward to the next opportunity to stop and check the roadmap.

### FOR INFORMATION

**Mikel Forcada** is Associate Professor in the Department of Software and Computing Systems at the Universitat d'Alacant (University of Alicante)

**Email:** [mlf@dlsi.ua.es](mailto:mlf@dlsi.ua.es)

**Web:** [www.dlsi.ua.es/~mlf/](http://www.dlsi.ua.es/~mlf/)

**Harold Somers** is Professor of Language Engineering in the Centre for Computational Linguistics at UMIST

**Email:** [Harold.Somers@umist.ac.uk](mailto:Harold.Somers@umist.ac.uk)

**Web:** [www.ccl.umist.ac.uk/staff/harold/](http://www.ccl.umist.ac.uk/staff/harold/)

**TMI 2002** (9th International Conference on Theoretical and Methodological Issues in Machine Translation) **web site:** [www.kecl.ntt.cqjp/events/tmi/](http://www.kecl.ntt.cqjp/events/tmi/)

(continued from page 7)

On further investigation, it turns out, not surprisingly, that there is a lot of structure among reviewers, committees, submissions, acceptances, and participation. Groups that are well represented in any of these tend to be well represented in the others as well. There is also quite a bit of structure over time. Groups that are well represented in previous meetings tend to be well represented in the future. Thus, reaching out to a market segment can reap benefits to the society for years to come.

A great way to reach out to a growth segment is with assignments (reviewing, positions on committees, etc): too often we treat these potential carrots as thankless chores and arm-twist the same people to do them again and again. In some cases, we even squander potential carrots by assigning them for life; recycling the carrots by handing them out to different people every few years would reach a larger market. We can also increase the number of carrots by maximising the number of reviewers, creating new posts, and so on. A not-so-great way to reach out to growth segments is blind reviewing. I have no objection to blind reviewing, but even if it is

effective (and I am not aware of data demonstrating that it is), blind reviewing is too passive. The society ought to market itself more aggressively than that – by identifying core segments in which it is already strong, and targeting each year a few segments for growth. This is not so much an issue of fairness, but a simple matter of what is in the best long-term self-interest of the organisation.

In summary, we ought to identify a few areas for growth and nurture them in two ways:

- **Effective use of carrots** (reviewing, positions on committees, etc.). Increase the number of carrots as much as possible to make sure that there are plenty of carrots available for reaching out to targeted growth areas.
- **Correction for reviewers' natural bias towards conservative precedents.** Accepted papers should have an advocate, a second, and no serious objections. Abstentions don't count. Make sure that papers in targeted areas are not rejected for bad reasons (e.g., abstentions, lack of precedents).

e

# Opinion

Opinion Column

*Kenneth Church, AT&T Labs Research, USA*

A couple of years ago there was a rather unpleasant debate in this column on affirmative action. Should organisations like ACL and COLING take a pro-active stand to make sure that certain groups are appropriately represented? I don't want to get into great questions such as what's right and what's wrong and what's fair and what's unfair. That is all much too complicated. And it depends on your perspective. Such questions are too divisive: I prefer to frame the debate in terms of simple market forces that we can all appreciate.

At the end of the day, the future of an organisation depends on growth. Even Wall Street, which is hardly known for its long-term vision, makes a big distinction between top-line revenue and bottom-line profits. You might think that money is money and that it wouldn't matter to a stock whether the company increased its revenue or decreased its cost. But Wall Street is willing to pay a premium for growth stocks which emphasise top-line revenue at the expense of short-term profits. Even Wall Street, which has little concern for what's right and what's fair – and not a lot of patience – appreciates that growth is important: we can't keep doing what we already know how to do.

Of course, growth is hard work. Old dogs don't like to learn new tricks. There are a lot of standard excuses for doing the same old thing: let me refer to these arguments as the conservative position. A standard conservative argument is quality: let the review process do whatever it does and then accept the papers that get the best reviews, and reject the others. It sounds simple and fair. But without offering a premium for growth (e.g., new sources of ideas, members, etc.), the conservative position ends up accepting the same old papers from the same old places. That's a death spiral. Old dogs eventually die.

The conservatives assume that the standard review process works pretty well. We all know better, of course. Having organised a number of conferences over the years, I've seen that reviewers take their job seriously and do the best they can (for the most part), but even so, the process is far from perfect. After a particularly boring paper, I am often asked why the programme committee acted as they did. Such questions are never easy; I hate to admit that we make mistakes, but we do. The review process is good for, at most, a few bits. It can separate the top half from the bottom half with reasonable accuracy, but it is asking too much of the process to reliably separate the top 10% from the next 10%. If the acceptance rate falls below 20% or so, you can be sure that there will be some mistakes.

A big problem with the review process is that reviewers are like lawyers. They love precedents. Reviewers don't

know what to do with papers that break new ground. The grades are more random for innovative papers than for incremental papers. All too often, one of the random grades comes out low and kills a good paper. The fundamental problem is that the standard reviewing process is inconsistent with the long-term needs of the society: reviewers are looking for easy precedents whereas the society needs a constant stream of new ideas from new people with new perspectives.

The Chair can and should correct for some of these biases by changing the rules. Instead of merely averaging the grades, it helps to accept only those papers that have an advocate, a second, and no serious opposition. I have found that this practice helps weed out incremental papers, which tend to have high average grades but few advocates. This practice also helps innovative papers where the low grade is often more of an abstention than a serious objection. (By serious objection, I am looking for a highly confident and clearly correct review that points out a major error.)

I like to think of running a conference in terms of marketing. If you segment the market (defined appropriately in terms of submissions, acceptances, reviewers, participants, or whatever) by almost any variable (e.g., geography, institution, topic, gender), the distribution tends to be skewed, perhaps even Zipfian. There are a few core segments that are well represented and many growth segments that are less well represented. The society ought to nurture the growth segments since the future depends on them.

I served on one programme committee where we did whatever we did, and then we checked how we had done in all the standard ways, and everything looked OK. And then we did a non-standard check, where we stratified countries by the number of submissions. To our horror, we realised that there was a huge effect. Countries that had submitted lots of papers had much higher acceptance rates than countries that had submitted just a few papers. You couldn't see the effect if you looked at each country one at a time, but if you grouped countries by the number of submissions, the pattern was all too clear.

*(continued on page 6)*

## FOR INFORMATION

**Ken Church** is Head of the Department of Data Mining at AT&T Labs Research in New Jersey

**Email:** [kwc@research.att.com](mailto:kwc@research.att.com)

**Web:** <http://www.research.att.com/~kwc>

**Spring  
2002**

**elsnet**  
.....

# SIGdial: Current Research Activities on Discourse and Dialogue in Japan


*Yasuharu Den, Chiba University, Japan*

SIG-SLUD (Spoken Language Understanding and Discourse Processing) is one of the biggest communities involved in discourse and dialogue research in Japan. It was formed in 1992 as a special interested group (SIG) of the Japanese Society for Artificial Intelligence and has a longer history than SIGdial. The current membership is approximately 140. The current president is Professor Akira Kurematsu of the University of Electro-communications, who is also a member of the Science Advisory Committee of SIGdial. SIG-SLUD has three meetings every year and, on average, twenty-five technical papers and four invited tutorials/lectures are presented annually.

SIG-SLUD and SIGdial have many points in common. One important issue common to the two communities is the development of shareable linguistic resources and standards for dialogue annotation. For this purpose, SIG-SLUD established a working group on discourse tagging in 1996 and its successor, a working group on corpus-based research on discourse and dialogue, in 1998. These working groups consist of tens of researchers not only from computational linguistics but also from various fields in the humanities. They have developed schemes, as well as annotation tools, for utterance segmentation, dialogue-act and exchange-structure annotation, and discourse segmentation for Japanese dialogues. The working groups have also made a great effort to annotate, using these schemes and tools, the dialogue corpora that have been collected at various universities and institutes belonging to the groups. The collection of annotated corpora amount to a running time of nearly 300 minutes, and will be published in mid-2002.

Another feature of SIG-SLUD that is shared with SIGdial is the encouragement of empirical methods, both in scientific and technological research. One of the three meetings each year specifically addresses corpus-based research. In particular, some members of the above-mentioned working groups present preliminary results based on the schemes and the corpora under development, which, together with feedback from a wide range of researchers, contributes to refining the schemes.

Although SIG-SLUD and SIGdial have much in common, SIG-SLUD also has some unique characteristics that SIGdial might not have. A distinctive feature of SIG-SLUD is the breadth of the research areas which it targets. In the past three years, approximately one third of the technical papers have been presented by researchers from fields other than computational linguistics or other computational areas. These have



*Yasuharu Den*



included phonetics, pragmatics, discourse analysis, conversational analysis, psycholinguistics, social psychology, and developmental psychology. Many of the invited tutorials/lectures have also been targeted at topics in social psychology, ethnomethodology, developmental psychology, and other humanities/social science fields related to communication. The interaction between the information sciences and the humanities/social sciences has made discourse and dialogue research activity in Japan extremely attractive. This multiplicity of disciplines has also been an advantage in the development of the dialogue annotation schemes, which are not restricted to small application domains.

Researchers in the information sciences are, of course, active in guiding the community. They notice the significance, in discourse and dialogue research, of advances in other areas and still keep track of their goals in developing applications. The most recent trend in this direction is the employment of VoiceXML as a vehicle for studying a general architecture for spoken dialogue systems. It enables us to implement various theories on discourse and dialogue developed so far, providing a solution to integrating different processing modules and to sharing discourse processing libraries among the community.

In summary, the discourse and dialogue research community in Japan is very active, manifesting its unique aspects as well as sharing much with the international community. I hope that in the near future a SIGdial workshop will be held in Japan, possibly in conjunction with ACL 2003 in Sapporo, and that fruitful interactions will take place between SIGdial and our community.

## FOR INFORMATION

**Yasuharu Den** is Associate Professor in the Faculty of Letters at Chiba University, Japan

**Email:** den@L.chiba-u.ac.jp

**Web:** cogsci.l.chiba-u.ac.jp/~den/

**SIGdial Website:** <http://www.sigdial.org/>





# New ELSNET Members

Membership Update

ELSNes would like to welcome new (and fairly new) members to ELSNET: Neurosoft, from Poland, who joined in January this year; Reuters (London), the SEMA Group (Madrid), and Universidad Nacional de Educación a Distancia (UNED, also in Madrid), who all joined during 2001.

A particularly warm welcome also goes out to Jose Maria Cavanillas of the SEMA Group, who recently joined the ELSNET Board.

ELSNes invited UNED and Neurosoft (page 10) to give us an idea of the type of work they do.

## The UNED group in Natural Language Processing

*M. Felisa Verdejo, UNED, Madrid*

The Spanish Open University (UNED for short) began its activities in October 1972 with the aim of making higher education accessible to anyone with the necessary aptitude and interest, regardless of formal qualifications. The UNED has proved an enormous success, and it is numerically one of the largest universities in Spain.

The UNED NLP group, comprising eleven members, is included in the Department of Lenguajes y Sistemas Informáticos (Computer Languages and Systems) with teaching activities in industrial engineering and computer science. It offers postgraduate courses on:

- Natural Language Processing
- Information Retrieval and NLP
- Design of Interactive Learning Environments
- Artificial Intelligence

These courses are also integrated within a Spanish inter-university doctoral programme in Cognitive Sciences.

The UNED NLP group started its research activities in 1991, and has been involved in basic and applied research concerning the design and evaluation of interactive and multilingual information retrieval systems, lexical knowledge acquisition and representation techniques, word sense disambiguation algorithms, natural language interfaces, and discourse modelling for teaching and learning environments.

The current interests of our research can be clustered around two broad topics:

### Multilingual, Interactive Information Retrieval

We are interested in the application of robust, scalable NLP techniques to improve interactive retrieval in realistic search situations, especially over multilingual collections and web sites. As this implies, we also have an interest in developing sound methodologies for the evaluation of these tasks.

Our most recent activities include:

- Development of the **Web Site Term Browser** – an interactive, multilingual search facility that exploits phrasal information extracted from portal-like domains to improve browsing and searching processes
- A system to support cross-language document selection based on a translation/summarising strategy
- Involvement in the organisation of **CLEF** – an evaluation campaign for multilingual information retrieval systems. Our participation includes the definition and coordination of an interactive cross-language retrieval track.

### Acquisition of Semantic Knowledge and Word Sense Disambiguation

Our research goal is the enrichment and improvement of (multilingual) large-scale lexical resources, mainly mining resources such as corpora, dictionaries, and the internet. In relation to this, we are also working on the development of unsupervised, robust word sense disambiguation techniques that have been tested in the recent SENSEVAL 2 competition.

### Research Results

Besides scientific publications (available from our web site), we have contributed to the creation of the following resources and products:

- the EuroWordNet semantic network – a multilingual lexical database with semantic relations between words in eight European languages, available through ELDA
- the Spanish SENSEVAL 2 test suite for the (comparative) evaluation of word sense disambiguation systems
- the CLEF test suites for the evaluation of multilingual information retrieval systems, comprising a comparable corpus of several gigabytes of news from agencies and newspapers in seven European languages, and a set of hundreds of queries with manual relevance assessments
- A phrase-based, multilingual site-searching facility (available at <http://rayuela.lsi.uned.es/wtb>)
- An interactive, web-based course on natural language processing and information retrieval (available at <http://rayuela.lsi.uned.es/~ircourse>).

### FOR INFORMATION

**M. Felisa Verdejo** is Professor in the Department of Computer Languages Systems at UNED

**E-mail:** felisa@lsi.uned.es

**UNED NLP Group:** <http://sensei.lsi.uned.es/NLP>

Spring  
2002

*elsnet*  
.....

# Neurosoft: Theory and Practice

*Cezary Grzegorz Dolega,  
Wrocław, Poland*



Neurosoft was established in December 1992 as a privately-owned company with headquarters in Wrocław (Poland). From the outset, Neurosoft has worked in close cooperation with the Polish scientific community especially with the Institute of Engineering Cybernetics of the Wrocław University of Technology and the Institute of Computer Science of the University of Wrocław. Our work in NLP technology is systematically verified by experts from the Polish Academy of Science (Institute of Computer Science in Warsaw) and Warsaw University (Institute of Polish Language).

The primary (and ultimate) goal of Neurosoft is to find practical applications for some of the theoretical ideas in the subject of Artificial Intelligence (AI). We believe, perhaps a little naively, that "the most practical thing is a good theory". The name "Neurosoft" comes from artificial neural networks, the most promising AI concept of the early '90s. The first two projects we were involved in concerned speech synthesis and recognition, and document processing (OCR, ICR), and were the continuation of research projects that had been conducted at the Institute of Engineering Cybernetics. As a result, in 1994, Neurosoft completed its first commercial product, **SynTalk<sup>®</sup>**, a text-to-speech (TTS) synthesis software for the Polish Language. Today *SynTalk* is the most popular TTS software in Poland (with nearly 60000 licences sold and over 5000 registered users) and runs on several platforms (MS Windows, PDA, Linux, QNX).

In comparison to speech synthesis, speech recognition and document analysis appeared to have far more complicated problems. Since 1992 Neurosoft has performed many research tasks relating to the recognition process. A unique neural network architecture has been constructed – **Spatiognitron** (spatiotemporal and hierarchical ANN based on the Time Delay Neural Network concept) – which performs segmentation-free recognition of a continuous signal, i.e., continuous speech (time domain) or continuous writing (spatial domain). Spatiognitron has been successfully applied to low quality printing recognition. However, the results of experiments indicated that the "single neural network" solution is less effective than the "multi-expert" method, which involves both the neural network and the structural recogniser. Building on this knowledge, at the beginning of 1997 Neurosoft launched its second commercial product, **Neurosoft BIP<sup>®</sup>**, which is a specialised document analysis and OCR engine.

Five years of intensive work on speech and writing recognition have emphasised to us the clear requirement for an advanced, high-level output-verification module in any

effective recognition tool (particularly for continuous speech recognition). Such a

module has to utilise any contextual information available: in general, this is lexical, syntactic, and semantic information. Moreover, there must be some kind of feedback from the output-verification module which changes the parameters of the low-level recognition stage. The application of NLP techniques seems to be the most natural and practical solution here. NLP methods are also very helpful in speech synthesis: the speech prosody generation algorithm (currently developed as a part of the *SynTalk 2.0* project) applies syntactic decomposition data to the initial output.

In 1994 Neurosoft began the construction of language processing tools to facilitate high-level verification. The first tool, a simple dictionary-based proofing module (which used a complete dictionary for Polish), produced too much ambiguity. For that reason, Neurosoft decided to build a more sophisticated mechanism involving syntactic and semantic analysis. In January 2000 the first version of a new commercial product, **Neurosoft Gram**, appeared on the market. This was primarily a typical morphological analyser dedicated to the Polish language, and capable of performing such tasks as text segmentation (sentence and word decomposition) and part-of-speech tagging. We discovered fairly quickly that such a morphological processor could be successfully integrated with full-text search engines, which nearly always have problems with processing Polish. The current version of Gram has been extended, with its new features helping developers to integrate it with search engines.

Neurosoft is still looking for new applications for its products. Over the last two years we have been developing a new system, **Neurosoft Lex**, which is a set of specialised tools for processing legal data, and whose main goal is to provide accurate and instantaneous information on legal regulations in Poland.

## FOR INFORMATION

**Cezary Dolega** is Head of the R&D Department at Neurosoft

**Email:** [cezar@neurosoft.pl](mailto:cezar@neurosoft.pl)

**Neurosoft Web Site:** <http://www.neurosoft.pl>

**Neurosoft Gram** demo (in Polish) is available online at <http://www.neurosoft.pl/gram>. Several Internet portals and Polish search engines currently use Gram (e.g., NetSprint.pl)

**Neurosoft Lex** is available online (in Polish) at <http://serwis-prawnyp1>, but a version for German Law and EU regulations is also being prepared

e

# Integrating Speech Technology in (Language) Learning: Review of InSTIL 2002

Workshop Review

**Steve Larocca**, *US Military Academy (USMA)*; **Kathleen Egan**, *DARPA*; **Fenfang Hwu**, *Bowling Green State University, Ohio*; and **Philippe Delcloque**, *Manchester Science Enterprise Centre*

The first InSTIL event on American soil took place on March 26-27 in Davis, near San Francisco, California, as a satellite event of the large USA Computer Assisted Language Learning (CALL) Conference, CALICO. The two-day programme included an introductory seminar and an advanced workshop on the integration of speech technologies in language learning. Many presenters who had previously been at the 1998 Speech Technology in Language Learning conference (STILL '98), Marholmen, Sweden and InSTIL 2000, Dundee, Scotland had made the trip to California, thus forming a link between two famous bridges – the Forth Road Bridge and the Golden Gate Bridge. One of the early speakers pointed out how fitting the location was for a group born in Edinburgh which continues to build a bridge between the CALL and Speech communities. The event was sponsored by BetterAccent.com, SRI International, and leading CALL publisher Swets & Zeitlinger.



Three early pioneers from the USA. From left to right: Farzad Ehsani, Kathleen Egan, and Steve Larocca

Presenters came from all parts of the globe; in particular, Europe, North America, and Japan. Universities represented included Gifu, Kyoto, Nancy, Nijmegen, San Francisco, Tokyo, UMIST, and West Chester. There were many participants from the USA diplomatic and defence establishment, including the USMA at West Point. A number of companies and commercial players were also represented; namely, BlueShoe Inc., Extempo Systems Inc., Ordinate Corp., SRI International, Sehda Inc., and Virage Inc.

The introductory seminar covered the major thematic areas in the field: automatic speech recognition (ASR); audio synthesis; audio-visual (A-V) synthesis including Talking Heads; visualisation; and virtual reality (VR) prototypes. The contents of the one-and-a-half day advanced workshop marked a clear progression from September 2000 in Dundee.

In the morning seminar, Kristin Precoda from SRI gave an insight into the differences between how machines and humans perceive language, in a paper entitled *ASR: A different way of hearing, or, why did it do that?*

Precoda provided a solid framework for understanding and designing language models that will make the most of the technology. The final contribution to the seminar featured the work of Delcloque, Hwu et al. on the *History of Speech Technology in Language Learning*, exhibited in Europe in 2001, which allowed participants to take stock of what had been achieved in the last 40 years.



Kristin Precoda from SRI (left) and Fenfang Hwu (right). Hwu was the chief researcher for last year's *History of STILL*

The opening keynote speaker in the advanced workshop, Jared Bernstein from Ordinate, kicked off with an explanation of the use of ASR in language testing. InSTIL Board member of the ISCA



Jared Bernstein giving his opening talk

constituency and early pioneer of ASR in CALL, Jared was also asked to make the traditional after-dinner speech at the InSTIL social event held at the end of the first day. His witty contribution covered the part played by the West Coast laboratories, researchers, and commercial entities in the development of speech recognition.

Pronunciation training strategies utilising ASR were addressed at both theoretical and practical levels by

Spring  
2002

elsnet  
.....



Precoda, Tsubota, Raux, Cavalli-Sforza, Jacome, Ehsani, and Koster, who all represented talented multi-authored teams of researchers. The early part of the workshop was devoted to recognition.

The growing strength of prosody research in CALL was demonstrated in papers on stress, rhythm, and intonation by Molholt, Hamel, Girard, and Minematsu, featuring the use of synthesis as well as segmental/suprasegmental visualisation. There were also excellent contributions on speech enhancement and manipulation, such as the work of Balci, Colotte, and Bonneau from Nancy, France.

The importance of quality feedback was discussed at length by several teams of presenters. Neri presented the work carried out in Nijmegen; Ishi et al spoke about their work in Tokyo, and Voce described work under way at UMIST in Manchester.

A relatively new area of work included the use of animated agents, avatars, and generally VR in CALL. Hayes-Roth from Extempo and Delcloque from UMIST discussed such developments. The other equally new area is that of content extraction using ASR: Tanaka et al. and Merlino & Egan, respectively, gave papers which covered research in this area in Japan and the USA.



*Philippe Delcloque*

The InSTIL workshop was once again a small, dedicated, and convivial forum for a focussed discussion on this expanding area of speech research. It allowed presenters and participants to have in-depth debates on the various speech recognition toolkits, on advances in synthesis, and other technologies which underpin the spoken-language interface. It was noticeable that issues of integration are at the forefront, with examples of reading tutors, distance learning, speech-to-speech translation devices, interview tools, and so on. Less frequently learnt languages are also

appearing, with encouraging research in Arabic, for instance. The importance of speech technology within the field of CALL can now no longer be underestimated. The main conference itself, CALICO, featured some excellent work in the field and an exposé on the relative merits of authoring systems which integrate ASR, such as Authorware, Wincalis, and the very complete Blueglas

InSTIL 2002 ended with a vision of the present and the future given by closing Keynote Speaker Philippe Delcloque, who also informed the audience that a special session on STiLL will be held at the forthcoming ICSLP conference in Colorado, September 16-20, 2002. Philippe also announced the publication of the InSTIL journal later this year and the location of the next two-day InSTIL Symposium which is sure to attract a growing number of presenters. Organised by InSTIL Board member Rodolfo Delmonte, this will take place in June or October 2004 in Venice, a place where more bridges can be found! It looks like some InSTILers are already including this in their busy diaries.

This short review cannot give full credit to all presenters and their co-authors. For full details of all the contributors, the reader is invited to visit the new InSTIL web site (see below).

#### FOR INFORMATION

**Colonel Steve LaRocca** is director of a research cell in the Department of Foreign Languages at the US Military Academy, West Point, USA

**Dr Fenfang Hwu** is an Assistant Professor of Spanish linguistics in the Department of Romance Languages at Bowling Green State University, Ohio, USA

**Kathleen Egan** is Principal Investigator in Speech Technologies at DARPA's Office of Advanced Information Technology in Washington, USA

**Dr Philippe Delcloque**, founder of InSTIL, is currently Senior Enterprise Academic at the Manchester Science Enterprise Centre, UK

Any comments or queries to Philippe Delcloque

**Email:** p.delcloque@msec.ac.uk

**Web:** <http://www.instil.org>

InSTIL is a Special Interest Group (SIG) associated with ISCA (the International Speech Communication Association), CALICO (the Computer Assisted Language Instruction Consortium), and EUROCALL (the European Association for Computer Assisted Language). It has been running since 1996 (and used to be known as CAPITAL).

**InSTIL web site:** [www.instil.org](http://www.instil.org)

**ISCA web site:** [www.isca-speech.org](http://www.isca-speech.org)

**CALICO web site:** [www.calico.org](http://www.calico.org)

**EUROCALL web site:** [www.eurocall.org](http://www.eurocall.org)

Summer School  
Reminder



### Reminder and New Deadlines

The 10th ELSNET European Summer School on Language and Speech Communication, whose topic is the **Evaluation and Assessment of Text and Speech Systems**, will be held in Odense, Denmark, July 15-26, 2002.

Deadlines for pre-registration and grant applications have been extended.

For details, visit <http://summerschool2002.nis.sdu.dk/>

# Speech Research and Technology in Kaunas University of Technology

Feature

*Algimantas Rudzionis, Kaunas University of Technology*



ELSNews invited Dr Algimantas Rudzionis, Head of the Speech Research Laboratory at Kaunas University of Technology, to give us an idea of the work being done there in Speech Technology. Dr Rudzionis has written us this report, which discusses their relevant work – both past and present.

Research in various fields of speech processing spans almost three decades in Lithuania, with different institutions specialising in different fields of speech technology. Whilst Vilnius University has specialised in Lithuanian text-to-speech synthesis, the Speech Research Laboratory of Kaunas University of Technology (KUT) has concentrated on speech recognition problems. Indeed, this was the largest research group involved in the Lithuanian language during the Soviet time.

## Alternatives of Phoneme Discrimination

Even today, after some years of research in the field, the performance of speech recognition algorithms barely approaches that of humans. At KUT we believe that automatic speech recognition can be improved by using advanced phoneme discrimination methods.

Most of the existing speech recognition algorithms make the assumption that the phoneme is best represented by the set of frames which could be mapped onto a given Hidden Markov Model (HMM) state. Unfortunately, the relationship between the HMM state and the phonetic unit is not clear. Over the years we have developed a number of methods that have led to improved overall results in speech recognition:

- the **dichotomy-based** method – this being a modification of the Fisher-discriminant method, combined with the subspace principle, in which the basic phonetic unit is considered as a given consonant-vowel cluster. Our experiments involved the analysis of the consonants 'M', 'N', 'V', and 'L' before (or after) the mostly contrastive vowels 'A', 'U', and 'I' in English word pairs: e.g., "might – night", "moon – noon", "meat – neat"
- the **projection-based** method (PRJ), which allows the automatically detected consonant-vowel

transition to be exploited more widely. PRJ was developed as a phonetically motivated alternative in phonetic recognition, in which only a small number of automatically determined stationary and transitional feature-vector pairs are taken into account. The number of feature-vectors is very close to the number of phonemes or significant acoustic events in the whole utterance

- the application of **statistical classifiers** (neural nets, Fisher-discriminant analysis, etc.) with high discrimination capabilities. In fact, we postulate that the local discrimination capabilities of the most popular *continuous-density HMM-with-Gaussian-mixtures* model should be similar to the capabilities of the k-means classifier. Our initial experiments with different classifiers used the dichotomy-based method, with later experiments introducing the **regularised discriminant analysis** (RDA), which uses singular value decomposition and regularisation of eigenvalues and eigenvectors. The single-layer perceptron has also been used for classification
- the **expansion of the feature set** for phoneme recognition. Our experiments showed improved results when the widely-used mel-scale cepstrum (MFCC) features were supplemented with features derived from recursive digital filters.

## The Projection Algorithm

We have found the results of our experiments on the projection algorithm particularly encouraging. In initial tests for speaker-dependent isolated speech recognition, the algorithm outperformed dynamic time warping (the error decreased about twofold) using the same phonetically segmented features. In later work, the PRJ algorithm based on continuous features was improved by introducing the averaging of several pronunciations of the same utterance and combining the cepstrum and the recursive filter features to form a single feature vector. The speaker-dependent isolated word recognition error decreased linearly with increasing number of the averaged utterances. Finally, the PRJ algorithm was adopted for fully phonetic speech recognition when word reference is represented by its phonetic transcription. In the best case 99.1% recognition accuracy was achieved using between six and eight clusters per phonetic event. Phonetic training is very important, so contextual environments should be covered as much as possible.

&gt;

Spring  
2002

elsnet  
.....





*The home of Kaunas University of Technology*

### The Dichotomy-based Method

The initial encouraging experiments on speaker-independent discrimination of consonants 'M', 'N', 'V', 'L' before vowels 'A', 'U', 'I' were performed over the period 1985–90, using high quality speech. It was observed that for the discrimination of consonants before open vowel 'A' to achieve discrimination accuracy of about 2%, it is sufficient to use just the stationary part of the consonant. Similar discrimination accuracy before vowel 'U' could be obtained by joining the stationary part of the consonant with the transitional consonant-to-vowel segment. In the most difficult case when a consonant precedes the closed vowel 'I' much more detailed description was needed.

More recently, similar experiments have been repeated using telephone quality speech, simultaneously comparing the k-means and RDA approaches. The recognition accuracy decreased from 21.6 % for k-means to 3.6 % for RDA in the open vowel 'A' context, and from 27.9 % to 11.4 % in the close vowel 'I' context.

### Incorporating these Improvements into Speech Recognition Algorithms

The improved phoneme discrimination methods that have been developed at KUT would be expected to improve the performance of speech recognition systems into which they could be incorporated. There are a number of ways in which this could be done: for example, the phonetic discriminants could serve as additional features for continuous-density HMMs in a similar way to delta (or delta-delta) features. Alternatively, the discrete phoneme discrimination outputs could be used as the basis for additional small discrete HMM word reference, and so on.

### Other Recent Work at KUT

In addition to our work on phoneme discrimination, there are other areas of research in which KUT has recently been involved:

- the collection of a Lithuanian speech database, called **LTDIGITS**. This contains digit names, control words, isolated "nasal-vowel" syllables, and the same syllables in continuous phrases
- the development and implementation of a **stenography system** for the Lithuanian parliament. The aim of this is to assist the stenography process by controlling the playback of pre-recorded speech simultaneously with text collection from the Word text editor. This system could be modified for use by journalists, crime investigators, foreign language studies, medical purposes, etc.

We are currently involved in the development of the following national projects:

- Lithuanian text-to-speech synthesis and recognition for disabled people (web browsing by voice, internet reading by voice, voice driven on/off switching of electrical devices)
- voice-based services for Lithuanian Telecom and mobile service providers (answering machines, interactive voice response systems, spoken language interfaces).

### FOR INFORMATION

**Dr Algimantas Rudzionis** is Head of the Speech Research Laboratory at Kaunas University of Technology in Lithuania. The group participates in the European COST 278 "Spoken Language Interaction in Telecommunication" programme.

**Email:** alrud@mmlab.ktu.lt

e



# Future Events

- May 27-June 2** *Workshops in association with LREC 2002*: Las Palmas, Spain.  
There are 19 highly relevant pre- and post-conference workshops covering many areas related to language resources and evaluation. For details, see "Workshops" link on main conference web site.  
URL: [www.lrec-conf.org/lrec2002/](http://www.lrec-conf.org/lrec2002/)
- May 29-31** *3rd International Conference on Language Resources and Evaluation (LREC 2002)*: Las Palmas, Spain.  
Email: [lrec@ilc.pi.cnr.it](mailto:lrec@ilc.pi.cnr.it) URL: [www.lrec-conf.org/lrec2002/](http://www.lrec-conf.org/lrec2002/)
- June 10-12** *The International Semantic Web Conference*: Sardinia, Italy.  
Email: [missikof@iasi.rm.cnr.it](mailto:missikof@iasi.rm.cnr.it) URL: [iswc.semanticweb.org/](http://iswc.semanticweb.org/)
- June 10-21** *The LOT (Netherlands Graduate School of Linguistics) Summer School 2002*: Nijmegen, The Netherlands.  
Email: [lot@let.uu.nl](mailto:lot@let.uu.nl) URL: [www.lot.let.uu.nl/](http://www.lot.let.uu.nl/)
- June 17-21** *ISCA Tutorial and Research Workshop (ITRW) on Multimodal Dialogue in Mobile Environments (IDS02)*: Kloster Irsee, Germany.  
Email: [laila@nis.sdu.dk](mailto:laila@nis.sdu.dk) URL: [www.sigmedia.org/ids02](http://www.sigmedia.org/ids02)
- June 23-July 6** *NATO Advanced Study Summer School on Dynamics of Speech Production and Perception*: Il Ciocco, Italy.  
Email: [asi2002@ebire.org](mailto:asi2002@ebire.org) URL: [www.ciocco.it/](http://www.ciocco.it/)
- June 24-25** *Indo-European Conference on Multilingual Communication Technologies*: Pune, India.  
Email: [iemct-contact@cdacindia.com](mailto:iemct-contact@cdacindia.com) URL: [www.cdacindia.com/iemct](http://www.cdacindia.com/iemct)
- June 28-29** *International CLASS Workshop on Natural, Intelligent, and Effective Interaction in Multimodal Dialogue Systems*: Copenhagen, Denmark.  
Email: [kuppevelt@imsuni-stuttgart.de](mailto:kuppevelt@imsuni-stuttgart.de) URL: [www.class-tech.org/events/NMI\\_workshop2.html](http://www.class-tech.org/events/NMI_workshop2.html)
- July 1-3** *The Second International Natural Language Generation Conference (INLG2002)*: Ramapo mountains (near New York City), USA.  
Email: [rambow@research.att.com](mailto:rambow@research.att.com) URL: [inlg02.cs.columbia.edu/](http://inlg02.cs.columbia.edu/)
- July 6-12** *Workshops in association with ACL '02*: Philadelphia, USA.  
There are a number of relevant pre- and post-conference workshops associated with ACL '02. For details, see "Conference Workshops" link on main conference web site.  
URL: [www.acl02.org](http://www.acl02.org)
- July 7-12** *40th Anniversary Meeting of the Association for Computational Linguistics (ACL '02)*: Philadelphia, USA.  
Email: [joshi@linc.cisupenn.edu](mailto:joshi@linc.cisupenn.edu) URL: [www.acl02.org](http://www.acl02.org)
- July 15-26** *10th ELSNET European Summer School on Language and Speech Communication – Evaluation and Assessment of Text and Speech Systems*: Odense, Denmark.  
Email: [hensen@nis.sdu.dk](mailto:hensen@nis.sdu.dk) URL: [summerschool2002.nis.sdu.dk/](http://summerschool2002.nis.sdu.dk/)
- July 22-23** *International Workshop on Computational Approaches to Collocations*: Vienna, Austria.  
Email: [colloc02@oefai.at](mailto:colloc02@oefai.at) URL: [www.ai.univie.ac.at/colloc02/index.html](http://www.ai.univie.ac.at/colloc02/index.html)
- July 23** *Workshop on Semantic Authoring, Annotation, & Knowledge Markup (SAAKM 2002)* (in conjunction with ECAI2002): Lyon, France.  
Email: [handschuh@aifbuni-karlsruhe.de](mailto:handschuh@aifbuni-karlsruhe.de) URL: [saakm2002.aifb.uni-karlsruhe.de](http://saakm2002.aifb.uni-karlsruhe.de)
- July 27-31** *The Fifth Teaching and Language Corpora Conference (TALC 2002)*: Bertinoro, Italy.  
Email: [talc5@sslmit.unibo.it](mailto:talc5@sslmit.unibo.it) URL: [www.sslmit.unibo.it/talc](http://www.sslmit.unibo.it/talc)

This is only a selection of events – see <http://www.elsnet.org/cgi-bin/elsnet/events.pl> for details of more events.

Spring  
2002

*elsnet*  
.....

## ELSNET

## Office

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

## Task Groups

*Training & Mobility*  
Gerrit Bloothoof, Utrecht University (NL)  
Koenraad de Smedt, University of Bergen (NO)

*Linguistic & Speech Resources*

Antonio Zampolli,  
Istituto di Linguistica Computazionale (I) and Ulrich Heid, Stuttgart University (D)

## Research

Niels Ole Bernsen, NIS Odense University (DK) and Joseph Mariani, LIMSI/CNRS (F)

## Executive Board

Steven Krauwer, Utrecht University (NL)  
Niels Ole Bernsen, NIS, Odense University (DK)  
Jean-Pierre Chanod, XEROX (F)  
Jose Maria Cavanillas, SEMA (E)  
Björn Granström, Royal Institute of Technology (S)  
Nikos Fakotakis, University of Patras (EL)  
Ulrich Heid, Stuttgart University (D)  
Denis Johnston, BT Adastral Park (UK)  
Joseph Mariani, LIMSI/CNRS (F)  
José M. Pardo, Polytechnic University of Madrid (E)  
Tony Rose, Reuters (UK)  
Geoffrey Sampson, University of Sussex (UK)  
Antonio Zampolli, University of Pisa (I)

## The ELSNET Participants

## Academic Sites

A Austrian Research Institute for Artificial Intelligence (ÖFAI)  
A Vienna University of Technology  
A University of Vienna  
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 Universität Stuttgart-IMS  
D Christian-Albrechts University, Kiel  
D Universität Hamburg  
D Institut für Angewandte Informationsforschung  
D German Research Center for Artificial Intelligence (DFKI)  
D Universitaet des Saarlandes  
D Universität Erlangen-Nürnberg -FORWISS  
D Ruhr-Universität Bochum  
D Universität des Saarlandes CS-AI  
DK Center for Sprogteknologi  
DK University of Southern Denmark  
DK Aalborg University  
E University of Granada  
E Universitat Autònoma de Barcelona  
E Universidad Nacional de Educación a Distancia (UNED)  
E Universidad Politécnica Madrid  
E Universidad Politécnica de Catalonia  
E Universidad Politécnica de Valencia  
EL Institute for Language & Speech Processing (ILSP), Athens  
EL National Centre for Scientific Research (NCSR) 'Demokritos', Athens  
EL University of Patras  
F LIMSI/CNRS, Orsay  
F LORIA, Nancy  
F Université Paul Sabatier (Toulouse III)  
F Université de Provence  
F Inst. National Polytechnique de Grenoble  
F IRISA/ENSSAT, Lannion  
GE Tbilisi State University, Centre on Language, Logic and Speech  
HU Lóránd Eötvös University  
HU Technical University of Budapest

I Consiglio Nazionale delle Ricerche  
I Università degli Studi di Pisa  
I IRST, Trento  
I Consorzio Pisa Ricerche  
I Fondazione Ugo Bordoni  
IRL Trinity College, University of Dublin  
IRL University College Dublin  
LT Inst. of Mathematics & Informatics  
NL Utrecht University  
NL University of Amsterdam (UvA)  
NL University of Nijmegen  
NL Foundation for Speech Technology  
NL Tilburg University  
NL University of Twente  
NL Eindhoven University of Technology  
NL Universiteit Leiden  
NL Netherlands Organization for Applied Scientific Research TNO  
NL University of Groningen  
NO University of Bergen  
NO Norwegian University of Science and Technology  
P University of Lisbon  
P New University of Lisbon  
P INESC ID, Lisbon  
PL Polish Academy of Sciences  
RO Romanian Academy  
RU Russian Academy of Sciences, Moscow  
S KTH (Royal Institute of Technology)  
S Linköping University  
UA IRTC UNESCO/IIP  
UK The Queen's University of Belfast  
UK Leeds University  
UK University of Sunderland  
UK University of Cambridge  
UK University of Sheffield  
UK University of York  
UK UMIST, Manchester  
UK University of Essex  
UK University of Edinburgh  
UK University of Sussex  
UK University of Brighton  
UK University College London  
UK University of Ulster  
UK University of Dundee

## Industrial Sites

D Acolada GmbH  
D Philips Research Laboratories

D IBM Deutschland  
D DaimlerChrysler AG  
D aspect Gesellschaft für Mensch-Maschine Kommunikation mbH  
D Grundig Professional Electronics GmbH  
D Novotech GmbH  
D Verlag Moritz Diesterweg GmbH  
D Sympalog Speech Technologies AG  
D Varetis Communications  
D Langenscheidt KG  
DK Tele Danmark  
E SchlumbergerSema sae  
E Telefonica I & D  
EL KNOWLEDGE S.A.  
F Memodata  
F Xerox Research Centre Europe  
F LINGA s.a.r.l.  
F Systran SA  
F VECSSYS  
F SCIPER  
F Aerospatiale  
F TGID  
FIN Kielilone Oy  
FIN Nokia Research Center  
HU MorphoLogic Ltd.  
I LOQUENDO  
I OLIVETTI RICERCA SCpA  
LV TILDE  
NL Compuleer  
NL Knowledge Concepts BV  
NL Sopheon NV  
NL IP Globalnet Nederland BV  
PL Neurosoft Sp. z.o.o.  
RU ANALIT Ltd  
RU Russicon Company  
S Sema Infodata  
S Telia Promotor AB  
UK Imagination Technologies plc  
UK Canon Research Centre Europe Ltd  
UK Sharp Laboratories of Europe Ltd  
UK BT Adastral Park  
UK ALPNET UK Limited  
UK 20/20 Speech Ltd  
UK Reuters Ltd  
UK Vocalis, Ltd.  
UK Hewlett-Packard Laboratories  
UK Logica Cambridge Ltd.  
UK SRI International

## What is ELSNET?

ELSNET is the European Network of Excellence 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 academic

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. E-mail sent to [elsnet-list@let.uu.nl](mailto: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](mailto: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](mailto:elsnet@elsnet.org)

**Web:** <http://www.elsnet.org>

e