

elsnews

The Newsletter of the European Network in Language and Speech

Autumn 1998

A European Masters in Language and Speech

Gerrit Bloothoof, Utrecht University



Developing new human-computer interfaces requires expertise in both Language and Speech, and education and training in the two areas makes for successful employees in the language industry. For ELSNET members these facts are so obvious they are almost clichés. But that does not mean they are reflected in the relevant curricula: at the moment few European university departments provide a balanced programme that encompasses both.

In response to this, ELSNET has for some time now promoted the idea of a Masters degree in Language and Speech which would set specific requirements and a standard for high-quality Language and Speech education across Europe. This goal is now close to being realized: a EuroMasters is currently being developed, and in the year 1999-2000 the first students will be able to register for it.

Background

Perceiving the need for a European Masters in Language and Speech, ELSNET organised a meeting on the topic in the spring of 1996; this resulted in an application for funding to develop such a degree. The application, which was submitted under the

European Union SOCRATES Programme, was successful, and the resulting project kicked off in September 1997. It is currently about halfway and has 10 partners (see information on the next page).

Developing a programme in Language and Speech requires reflection on its structure, its content and its implementation. In the first year of the project, a model for the degree was developed and its content defined. In 1998-99, the project's focus will be on implementation and publicity; and in 1999-2000, the first students will be able to study within the new framework.

The model

Members of the project felt that the European dimension should be interpreted in such a way that in principle all European academic institutions and all (suitably qualified) interested students should be in the position to take part. But there are significant differences between current European systems for higher education, and in many European countries a Masters degree is not a legal degree as such; this is obviously a major obstacle to a single degree at European level.

In order to overcome this problem, the project has involved the main European organisations in the field of Language and Speech, ESCA (the European Speech Communication Association) and EACL (the European Chapter of the Association for Computational Linguistics), and given the EuroMasters a structure that is innovative and flexible.

Today's students in Language and Speech will be tomorrow's researchers, developers and engineers, who determine the future of the area. Education and Training are crucial, it's as simple as that. Providing a standard and an assurance of quality for integrated programmes in Language and Speech at European level is one of the main priorities in the area of training. We expect that the EuroMasters, which is currently being developed and which will accept its first student in the year 1999-2000, will be a major landmark in this respect. Gerrit Bloothoof fills you in on the details. But education in NL & S encompasses many other formats as well, and we highlight several of them in this issue.

Reactions? Send them to Mimo Caenepeel (elsnews@let.uu.nl).

Biting the Bullet	
Johan Myking	3
Long Distance Learning	
LE Training Showcase	4
The Future of the Humanities in the Digital Age	
Daniel Apollon	
Anders Eriksson	6
Life is a sequence of syllables	
ELSNET Summer School	8
ELSNET News	
New EB member	9
ELSNET history	
Antonio Zampolli	10
Smoke-filled rooms	
EB minutes	12
Or Else	
Yorick Wilks	14
NL evaluation	
Bente Maegaard	15
GATE	
Hamish Cunningham	16
Book reviews	17
Announcements	18
Future events	19

Nov 1998

elsnet


The model looks as follows. The degree will not replace any legal degree students may obtain from their university; instead, it will be defined by contents only, and implemented through a certification procedure. In other words, the European Masters will be a certificate, awarded by ESCA and EAEL to students who have fulfilled the conditions set by the Masters, which will enrich the student's legal national degree by giving it a 'quality' stamp. We anticipate that such an assurance of quality will be respected and valued in academic and industrial environments worldwide.

How the degree is implemented in practice will depend on the degree structure of each individual country. In the UK, for instance, Masters (MSc) programmes are well-established, and our proposals will fit in well with existing courses. In countries like Germany or the Netherlands, on the other hand, there is no legal Masters degree, but the contents of the EuroMasters will be part of studies leading to a Magister or Doctorandus title. In many cases, departments already offer courses that correspond to the content of the European Masters; but they usually fall under different 'study paths' (NLP and Phonetics/Speech Technology). We hope that such departments will make the effort to develop a new direction of study which combines Speech and Language.

Content

The contents of the Masters programme will focus on the following areas:

- Theoretical Linguistics
- Phonetics and Phonology
- Human Language Processing
- Natural Language Processing
- Speech Signal Processing
- Statistical Pattern Classification
- Language Engineering Applications



EAEL

Overall structure

The Masters will consist of at least nine months of course work, and at least three months of project work (preferably a traineeship in industry). A total of at least three months should be spent abroad (whether on courses or on project work). In addition to this all students on the scheme will gather at a week-long Easter School, which will offer tutorials on the latest developments in the field. As a final requirement, students will have to submit a concluding thesis based on their project work, which will be published as a report on the electronic student journal WEB-SLS.

Joining the scheme

The scheme as a whole will be managed by an accreditation board, appointed by ESCA and EAEL. This board will judge applications from departments wishing to join the scheme. Relevant departments will need to

- show how they will cover the contents of the European Masters, through existing courses and/or courses in cooperation and exchange with other departments in Europe, possible complemented by open and distance learning courses;
- give an overview of the contents of these courses, teaching methods, books, and the number of ECTS [European Credit Transfer System] credits related to each course, and submit recent exercises, exam questions, and project subjects to illustrate the level of the courses;
- describe what measures they have taken, or will take, to support student exchange; and
- indicate which examination procedures will be used and in which legal degree the Masters will be embedded.

Once the department's proposal has been accepted by the accreditation board, the responsibility for certification of a student is transferred to the local examination board. This board is then entitled to award the certificate on behalf of ESCA and EAEL.

FOR INFORMATION

The current project group consists of the universities of Edinburgh, Sheffield, Essex, Aalborg, Utrecht, Saarbrücken, Lausanne/IDIAP, Lisbon, Patras and Athens.

For more information on the EuroMasters, please contact

Gerrit Bloothoof

Email: Gerrit.Bloothoof@let.uu.nl

Web: www.cstr.ed.ac.uk/EuroMasters

ELSNNews is published at the Centre for Cognitive Science, University of Edinburgh using Aldus Pagemaker™. It is printed on recycled paper by Lutton Press, Ltd.

Editor: Mimo Caenepeel
Editorial team: Mariken Broekhoven, Steven Krauwer and Geoffrey Sampson.

ISSN 1350-990X

© ELSNET 1998.

FOR INFORMATION
Contributions to ELSNNews, and address corrections, should be sent to:

elsnews@let.uu.nl

Tel: +44 131 650 4594

Fax: +44 131 650 6626

Material for the next issue is due:

15 December 1998

The theme of the next issue of ELSNNews will be

Spoken Language Dialogue Systems

Don't be a passive reader! Send us your tips, ideas, questions or articles: elsnews@let.uu.nl (before 15 December)

Information management: biting the bullet

Johan Myking, University of Bergen

ELSNET bullet courses are short intensive courses on a specific topic. The second one, which took place in Bergen (Norway) this September, focused on the role of computer-based terminology management systems in translational information management.

Johan Myking, one of the organisers, looks back on the event.

The organisation of a course like this one is no small feat. Targeting and reaching the right audience is one of the first and most important tasks, and in our case this involved, amongst other things, mailing out 700 leaflets and 300 emails. All things considered, email and the web appear to be the most efficient media for publicity, although some participants apparently were not aware of the website — even after the course!

In all, the results seem to have matched the expectations, both for the organisers and for the participants. The course attracted 17 participants from eight countries, with Norway, for obvious reasons, most heavily represented. Most participants expected an introductory “crash” course - and they certainly got a crash during the ten hours of Day 2! Having said that, there were very few complaints about “hard work” on the evaluation forms handed out after the course.

Course topics had been selected to cover a variety of practical aspects of terminology, and the lectures were well-received. The set of (very competent) invited speakers included Klaus-Dirk Schmitz (on terminology management), Gerd Engel and Bertha Toft (terminology and machine-translation), Chris Cox (quality management) and Øivin Andersen (conceptology and documentation). There were a number of demos as well. All these topics were considered relevant; but my impression was that *machine-translation outcome* might have been assigned a more prominent place in the programme.

After the course, participants were invited to give suggestions and recommendations for future courses, and their feedback was clearly in accordance with current tendencies in the field of Terminology. *Corpus*

creation, the interaction between text and terminology, and automatic term extraction methods are topics that are likely to attract the interest of a wide range of professionals. The course's emphasis on technology and practical applications, rather than on theory for its own sake, also seems to have been popular and successful.

I noticed to my satisfaction that the course attracted people who were quite unfamiliar with the field of terminology beforehand, and thus did not belong to “the family”. This was how I myself had interpreted the very idea of a “bullet” course. But the course also attracted some linguists, and even some practitioners of computational terminography. Luckily, this mixture of backgrounds was not a problem. Some participants may perhaps have found the level too elementary at times, or some topics of secondary relevance; but the discussions — both on a formal and informal level — and the variety of topics ensured an exchange of knowledge that had something to offer to everyone. I don't know whether this is in accordance with the purpose of the bullet courses, but it certainly fits in well with some essential principles of Terminology, such as interdisciplinarity and mutual cooperation across professional borders.



Finally, one non-scientific but crucial point to be considered for future courses. The organiser is never to be blamed for good weather, but we should not equate “bullets” with long hours more akin to “shell-bombing”. So a tip for future organisers: make sure the conference room has windows, and the air-conditioning works!

Johan Myking is Associate Professor at the Department of Scandinavian Languages and Literature and the HIT Centre, University of Bergen.
Email: Johan.Myking@nor.uib.no
Web: www.hit.uib.no/



An unusual number of L&S people was spotted in Bergen this summer

Nov 1998

elsnet
.....

Long Distance Learning through interactive tutorials on the Internet

Interactive Learning is in vogue these days, but the concept seems to be particularly relevant for education in the Language and Speech area. Because the area is multidisciplinary to an unusual degree, students have to acquire expertise in a wide range of topics which cross the boundaries of traditional disciplines; this is often a problem within current departmental structures. Moreover, practical skills in listening, analysis and performance are crucial, particularly for researchers in the Speech area; but acquiring such skills typically involves a high amount of staff/student interaction, which is hard to achieve at a time when there is pressure throughout European universities to increase efficiency and decrease the use of scarce staff resources.

Interactive tutorials which are accessible on the Internet seem to offer a solution to these problems: they give students the means to draw on varied expertise throughout the community, and provide training in practical skills without the need for continuous supervision. In addition to this, they can increase the efficiency of self-study, facilitate communication between "virtual" communities of teachers or students working in a particular sub-discipline, enhance the quality of teaching materials, and reduce the effort involved in creating and duplicating them. Again, some of these advantages are particularly relevant for the Language and Speech area; for instance, the Internet allows easier access to linguistic resources and native speakers of various human languages, regardless of geographical location.

Pilot 1: Models of Speech Perception

This project is developing a web-based tutorial on the four models of speech perception which are regarded as fundamental in the development of Phonetic Sciences, namely:

- Motor Theory (A.M. Liberman and colleagues)
- Analysis by Synthesis (K.H. Stevens)
- Quantal Nature of Speech (K.H. Stevens)
- Hyper and Hypoarticulation (B. Lindblom)

The aim is to provide students with a comprehensive overview of these models, including the hypothesis and theoretical foundations on which they are based; illustrations, experiments, and findings which support them; and problems and criticisms.

The tutorial will be organised in hypertext format. It will consist of a text-based course, with links to non-textual materials such as small experiments, sounds samples, and animated figures.

The contractor for this project is **Francesco Cutugno** (cutugno@cds.cised.unina.it) of CIRASS, Università di Napoli Federico II.
Collaborator: **Cécile Fougeron** (ILPGA, Université Paris III)

In short, interactive tutorials appear to be an excellent way of making wide and good use of specialist knowledge, and they are likely to become important components in courses throughout Europe. ELSNET wants to encourage the development of such tutorials in the area of Language and Speech, and to assess their strengths and weaknesses. In February of this year the ELSNET board approved support for **six pilot projects** in the area, which are now well underway. The tutorials developed by the projects, which are very different in character and scope, will function as test cases to determine the strong and weak points of interactive learning. ELSNET's support will stimulate their development, and make them accessible to a wide audience for both testing and use.

The monitoring and evaluation of the projects will be done by members of the Socrates Thematic Networks' **ACO*HUM** (Advanced Computing in the Humanities) and **Speech Communication Sciences**.

Pilot 2: The Linear Predictive Vocoder

The tutorial developed in this project will interactively explain what a Linear Predictive Vocoder does. The user speaks into a microphone, and his or her voice is digitised and stored in the computer. Then the voice components — the fundamental frequency and the vocal tract parameters — are computed and depicted, and fed into the synthesis part of the vocoder. And this generates a synthesised speech signal, which the user can replay and compare with the original speech signal.

The main advantage of the tutorial will be its interactive nature. It will, for example, enable the user to manipulate the fundamental frequency contour, the number of prediction coefficients or the signal energy, and to listen to the result of these manipulations. Although the LPC vocoder is primarily a coding scheme, it can also be used as a model of human speech production. It makes it easy for the student to observe that human speech is a concatenation of the vocal chord signal (represented by the fundamental frequency signal) and the resonance characteristics of the mouth and nose cavity. And it is not only instructive but also exciting to study both the visual and the audible variations of the fundamental frequency, and the way they vary with different speakers and different emotions.

The contractor for this project is **Klaus Fellbaum** (fellbaum@naxos.kt.tu-cottbus.de) of the Lehrstuhl Kommunikationstechnik, Cottbus Technical University, Germany.

Pilot 3: Interactive demonstrations in speech and hearing

In recent years, collections of auditory demonstrations have appeared on CD. Now, with the advent of easy-to-build interfaces and powerful signal processing/display languages, it has become possible to create more interactive forms of auditory demonstration, and that is what this project is doing.

Virtually all auditory phenomena can be appreciated with greater immediacy if the user is able to control the main parameters underlying the effect, so this is an area that is highly suited to user-guided exploration. Some examples of such phenomena: the beating of pairs of tones close in frequency, explored as a function of the tones' frequency separation; the pitch of a complex as a function of the harmonics selected for inclusion; the effect on the short-term speech spectrum of the placement and duration of a time-domain window; the streaming of double vowels as a function of F0 difference; and the detectability of a tone in unmodulated or comodulated noise.

The contractor for this project is **Martin Cooke** (M.Cooke@dcs.shef.ac.uk) of the Computer Science Department at Sheffield University.

Pilot 4: Statistical Natural Language Processing

This project is developing a web-based tutorial on statistical natural language processing. It will consist of a well-structured core course on the main elements of statistical NLP; and more loosely structured extensions into various subareas of statistical NLP, mostly organised as projects, with tools and corpora (or statistics from corpora) being supplied as much as possible.

Besides tools that will be used to convert material into HTML, the project will make available statistical NLP tools that can be used by students on projects. These tools will include various programs for deriving statistics from corpora, and some statistical taggers and parsers.

The contractor for this project is **Joakim Nivre** (nivre@ling.gu.se) of the Department of Linguistics, Göteborgs Universitet.
Collaborator: **Brigitte Krenn** (Universität Saarbrücken).

FOR INFORMATION

For more information on the ELSNET Open Distance Learning projects, please contact **Gerrit Bloothoof** (Gerrit.Bloothoof@let.uu.nl)

Pilot 5: Information Retrieval (mono- and multilingual) with Natural Language Processing techniques

This project aims to orientate students in the area of information retrieval, and provide them with the knowledge and skills required to integrate language and software technologies when building an application.

The course is designed specifically for distance learners, who will work collaboratively, in small groups, on a joint project. These projects will involve a variety of tools and skills, such as

- testing the capabilities and performance of different Internet searchers with selected text queries
- expanding queries manually using WordNet, and comparing results, with the same Internet searchers.
- using the SMART system on PC (Linux) and other publicly available software (stemmers, etc), and performing different tests on their impact on IR precision and recall.

The contractor for this project is **Felisa Verdejo** (felisa@ieec.uned.es) of the Universidad Nacional de Educación a Distancia, Madrid.

Pilot 6: NLP courseware for web operation

This project has three related goals: to develop *web tools* for the efficient representation of linguistic structures as pedagogical materials; to provide a *platform* for practical exercises in grammar writing, implemented by a web interface to a parser; and to develop an introductory NLP *course* based on the modules defined above, thereby validating the approach.

The envisaged outcome will consist of:

- generic *JAVA modules* for enriching web courses in many areas of linguistics/NLP. This will not be a one-shot result, but a set of generic pedagogical graphics tools which will be applicable in many subsequent web courses.
- a *parser* with a web interface, suitable for integration into courses. The mode of presentation will initially (i.e. at the start of the student's course) be web-based, leading to independent use of the downloadable Windows version.
- a prototype introductory NLP *course*, which will appeal to a very wide audience.

The contractor for this project is **Koenraad de Smedt** (desmedt@uib.no) of the Institutt for lingvistikk og litteraturvitenskap, Universitetet i Bergen.

Collaborators:

William J Black (Centre for Computational Linguistics, UMIST), **Henk Schotel** (University of Nijmegen), and **Andy Way** (Dublin City University)

Nov 1998

elsnet
.....

Integrating information technologies and humanistic disciplines

*Interview with Daniel Apollon, University of Bergen,
by Yvonne Robberstad Bonète, University of Bergen*



Daniel Apollon

Nobody questions the role and relevance of Advanced Computing for research and teaching activities in areas such as Computational Linguistics and Phonetics. But for humanistic disciplines the advantages are perhaps less obvious. "Why should we bother?" asked one of the speakers at the recent conference on The Future of the Humanities in the Digital Age, which was held in Bergen, Norway this September. Can the humanities benefit from emerging information and communication technologies? Do different disciplines within the humanities share common developments that can be supported by information technology? And how can the digital age adapt to the particular needs of the humanities? .

Before the conference Daniel Apollon, one of the organisers, talked to Yvonne Robberstad Bonète about the challenges he sees for the humanisation of information technologies, and the role humanistic disciplines could play in the information society. What follows is an edited extract of that conversation. On the next page, Anders Eriksson looks back on the conference itself.

FOR INFORMATION

Daniel Apollon is the coordinator of EUROLITERATURE, a SOCRATES programme on Global communication and the future of literary studies. He was co-organiser of the conference on The Future of the Humanities in the Digital Age.

Email: apollon@psych.uib.no

Web: www.euroliterature.uib.no/partners.html

For more on EUROLITERATURE, see

www.futurehum.uib.no/

Robberstad: As the coordinator of EUROLITERATURE and one of the organisers of the conference on *The Future of the Humanities in the Digital Age*, could you say a few words about the main topics of the conference?

Apollon: Information technologies have changed profoundly during the last three decades. Some of these changes, such as the explosion of global communication through the Internet, have attracted a lot of attention. Others, of a more cultural and social character, have been less visible to the general public. The last five years have seen the spectacular, but still poorly understood, spread of information technology to cultural and educational sectors. If we look at the total volume of applications of information technologies, we see that an increasing number of cultural and public communication applications with an explicit human-meaning content are becoming increasingly central.

Although traditional applications such as process control (controlling power plants, managing traffic lights) and record keeping (employee data bases) are still vital and indisputably offer formidable challenges to computer scientists, the future development of information technology rests on the integration of symbolic, human-like communication.

In Europe, however, departments dealing with information technology are still concentrating their research and educational efforts on more traditional "computer science" activities. In the meantime, traditional IT, media and telecommunication are converging and merging into a global all-encompassing super-medium, which will shape future human communication and reveal new forms of expression.

The humanisation of information technologies — the development of ICT towards more human-like, intelligent systems accessible to and manageable by a general public of all ages — raises formidable challenges. A few of these challenges are purely technical and related to speed and reliability of software and hardware. The real challenges, however, are of a non-technical nature, and concern the integration of human symbolic communication (such as natural language interfaces), new social patterns (e.g. virtual communities) and new cognitive models (e.g. exploratory learning) into the new convergent supermedium. It is already clear that the humanisation of information technologies will change the face of education, publishing, academic communication and cultural expression.

Future information technologies will not take off without a significant contribution from humanistic disciplines. At the same time, these disciplines can no longer afford to limit their field of reflection and their applications to traditional media such as books.

— why we should bother

Conference report

Anders Erikson, University of Umeå

There are few initiatives in Europe which link educational development, especially in higher education, with the future of the humanities in the information society. The conference *The Future of the Humanities in the Digital Age*, held in Bergen, Norway at the end of September was a notable exception, and it attracted considerable attention.

At the heart of the conference were issues such as

- the integration of converging information and communication technologies and humanistic disciplines;
- common elements in the evolution of humanistic disciplines within the information society; and
- linking educational development, and particularly recent developments in Open and Distance Learning, with the particular needs of the Humanities.

While contributions covered a wide variety of topics — ranging from *Art on the Internet to Language engineering for translation curricula* — they all addressed the use of information technology in research and teaching in the humanities.

The extent to which IT is used varies a lot between disciplines. The speech sciences currently seem to be the only field where advanced computing tools are used to

any appreciable extent; but the digital age is here, and it is making its presence known throughout the humanities, just like in every other field. Is this good or bad or just inevitable? Or, to cite the somewhat provocative question used as the title of **Christian-Emil Ore's** keynote speech at the conference: *Advanced computing in the humanities: Why should one bother?* It seems obvious that computers are used to some extent because they are there and because using computers is fashionable, a sign of being in touch with the times; but one should perhaps take warning from studies which have shown that when computers were first introduced in small businesses in America, productivity decreased rather than increased as a result.

The use of computers is a good thing only if it meets a need and allows you to do the things you want to do, better or faster. This is not always true in humanities departments, and the result, as reported by many of the participants in this conference, is that many teachers and researchers in the humanities meet the introduction of computers with outright hostility. If we want to be guided by the specific needs of a given subject, rather than leaving the development of computing to trigger-happy enthusiasts, then a thorough, and very concrete, discussion of what Information Technology can do for the various subjects in the humanities is needed. Computational linguistics and phonetics, where advanced



Anders Eriksson

computing tools are widely used, are perhaps the fields where, at present, the use of computers is most clearly the response to the specific requirements of the subjects; and as a consequence hostile attitudes towards computers are rare or absent. Perhaps there is a lesson to be learnt here?

Apart from the many stimulating talks and workshops there was, of course, also a social programme, including a reception by the vice-mayor of Bergen, a banquet in a hilltop restaurant, and several art and music events, offering a taste of what Bergen has to offer as the elected cultural capital of Europe for the year 2000.

FOR INFORMATION

The conference took place from 25-28 September 1998. It was a joint initiative of the SOCRATES/ERASMUS thematic network on Advanced Computing in the Humanities (ACO*HUM) and the SOCRATES/ODL project EURO-LITERATURE.

For more information about the conference, abstracts and recordings of some of the speeches, see www.futurehum.uib.no

Anders Eriksson is Senior Lecturer at the Department of Phonetics, University of Umeå, Sweden.

Email: anderse@ling.umu.se,

Web:

www.ling.umu.se:80/~anderse



Looking to the future: Bergen in September

Nov 1998

elsnet
.....

“Life is a sequence of syllables”



During July of this year the Universitat Politècnica de Catalunya (UPC) and the Universitat Autònoma de Barcelona (UAB) jointly organised the 6th European Summer School on Language and Speech Communication. This year's topic was *Robustness: Real Life Applications in Language and Speech*. Over 120 people (half of them from Eastern Europe) applied for a place at the Summer School. Unfortunately the optimum number of participants was only around 60, and so the organising committee had to make some hard decisions. 68 people eventually took part (10 academic staff, 41 (mostly PhD) students and 17 participants from industry).

What follows are four takes on the Summer School from different vantage points..



The academic's perspective

Bojan Petek, University of Ljubljana

Traditionally, ELSNET Summer Schools focus on the integration of speech and natural language processing, with emphasis on issues considered to be important for both research areas. This year's topic, robustness, is of fundamental importance to speech perception, and we strive to understand how the human auditory system achieves this important property. In speech technology we explore techniques or models that show promise in making automatic speech recognition systems robust. And Natural Language Processing components have to cope with transcriptions of spontaneous speech or dialogues that differ significantly from the ones obtained from written language. So robustness becomes a central theme of investigation for any serious language or speech application envisaged to solve a real world task.



This year's summer school programme contained a wide variety of topics under the robustness umbrella, all very relevant and all presented by leading experts in the field. During the second week **Steven Greenberg** presented some excellent and easy-to-follow plenary lectures on *The Processing of Spoken Language in the Real World*. He discussed several fundamental issues, for instance the mechanisms that pertain to robustness of the auditory system, the canonical properties of spoken language, and computational approaches to the study and comprehension of spoken language. Very illuminating, too, were his extrapolations on future speech understanding systems, which will exhibit structural properties significantly different from the current state-of-the-art HMM-based approaches. In particular, insights obtained from the **Switchboard** corpus suggest that syllables might become the units to model in systems for English spoken language understanding.

Last but not least, it was good to be able to meet colleagues and new friends in Catalunya. And equally enjoyable was a 2-day motorcycling tour to Cadaques (Costa Brava) and Vielha (Val d'Aran)— many thanks again to Ignasi for his brilliant guidance in planning this memorable touristic experience.



Bojan Petek (bojan.petek@uni-lj.si) is Assistant Professor at the Faculty of Natural Sciences and Engineering Voice, University of Ljubljana, Slovenia.

e

The 6th European Summer School on Language and Speech communication

The brain researcher's perspective

Yury Shtyrov, University of Helsinki

My background is in brain research, and in spite of my overall literacy in speech science I feared that courses at the ELSNET Summer School might be hard, even too hard, for me to follow. But surprisingly most lecturers, both in the plenary sessions and the speech recognition course slots, managed to convey the complex material they had prepared in a clear and absorbing fashion, without overloading it with too many technical details. No surprises followed this, though: the problem of speech recognition is still way too far from being solved, and we heard plenty of examples of tries and failures from brilliant teachers and their equally brilliant students.

The most unpleasant part for me personally was certainly my own presentation: I was afraid that nobody with a computer science or linguistic background would be able to understand what brain research could possibly have to do with speech. The result, however, was overwhelming: my report on speech processing in the brain elicited a lot of interest and numerous questions and comments, and people kept asking me questions for days after. The approach behind most of them may have seemed naive — “how can we copy the brain to build successful speech recognizers?”; but it was this approach that seemed to dominate the second week's courses, when **Steven Greenberg** and **Martin Cooke** tried to persuade the audience that it is only by studying and modeling mechanisms of sound and speech perception in humans that speech science can reach its ultimate goal. This was not only interesting but seemed very promising too (at least, to me), and it looks like an emerging “hot topic” in current speech science.

Yury Shtyrov (shtyrov@helsinki.fi) is a researcher at the Cognitive Brain Research Unit, Department of Psychology, University of Helsinki

The industrial perspective

Alessandro Maro, CSELT

“We need robust applications” — that's the strongest message I got attending the first week's lectures at the ELSNET 1998 Summer School. Both **Jean Claude Junqua** and **Daniel Tapias** stressed this point: real life applications do need robust recognition in order to achieve satisfactory performances. And that is dramatically true. Robustness is no longer an academic question, it is a real need. The market is now demanding new solutions and services for computer telephony and the car domain. Most current commercial systems are not yet able to handle phenomena such as speaker and channel variability, background speech or noise, and so on. In real applications, when the testing conditions differ from the training, word accuracy degrades, and so does customer satisfaction. Companies are sending this message to research groups: they pay for the research and they want to see that investment pay off. To

achieve that, researchers must supply efficient and robust solutions to satisfy market demands as soon as possible.

In my opinion one of the aims of the ELSNET Summer School was to warn young researchers that companies want to meet these demands now. I think they got the message.

Alessandro Maro (alessandro.maro@cse.lt.it) works at CSELT (Centro Studi E Laboratori Telecomunicazioni S.p.A.), Turin, Italy

The Speech perspective

Daniela Raddino, Sony International Europe

Robustness is one of the most important characteristics of a speech system if the system is to function well in different contexts: its success or failure may depend on it. Users need to be able to interact with the system in a spontaneous rather than a mechanical way. So the topic of the ELSNET summer school is an important and timely one.

In my view, the ELSNET summer school achieved a number of things. First, it provided an opportunity for the exchange of information between researchers from different companies and institutes working on speech recognition. Second, the lectures treated both new research topics and recent development of basic research items; that kind of balance is important. And finally, various perspectives on robustness were offered, including that of natural language processing, which was treated thoroughly.

In terms of individual courses, it was interesting to investigate new ways to treat the signal in the course *The Processing of the Spoken Language in the Real World* by **Steve Greenberg** and in *Robust Speech Processing in the Auditory System* by **Martin Cooke**. And *Robust Speech Recognition for unknown Compensation*, taught by **Jean-Claude Junqua**, offered insights fundamental to understanding the development of basic research topics such as speech recognition in noisy and variable environments, and speaker variability.

Daniela Raddino (raddino@sony.de) is a Development Engineer at the European Research and Development Stuttgart/ Stuttgart Technology Center, Sony International Europe.

FOR INFORMATION

The Summer School web pages are at
gps-tcs.upc.es/veu/ess98

Nov 1998

elsnet
.....

ELSNET News

Steven Krauwer, Utrecht University

Another monument of ancient ELSNET times has left us: **Ewan Klein**, who was one of the founding fathers of ELSNET, has decided to step down from the ELSNET Executive Board. Ewan was the Coordinator of the network from the very beginning of ELSNET until the end of 1994. He started with nothing but a list of people and institutions which might be interested in joining, and when I took over in 1995, he left a flourishing network with over a hundred nodes, full of activity and full of completed, ongoing, and planned actions.

If you ask me how exactly he did it, I must confess that I don't know. I didn't join until the real hard work had been done by Ewan and his colleagues on the ELSNET Board. My earliest contacts with Ewan were, as far as I remember, my requests for financial support for EACL 1993, which he turned down, very politely and on very well-motivated grounds. I am sure that many of you have gone through the same experience. I think that this was (and still is) his secret weapon: once he had got his arguments right, he could present them in a pleasant, almost charming way, so that you could never really feel offended, even if you didn't agree.


After Ewan resigned as ELSNET's coordinator he spent a year in the US, and then came back to the ELSNET Board, where we could

all benefit from his experience and his creative ideas.

Thanks, Ewan, for what you did to get ELSNET off the ground, and to keep it running for such a long time!

Ewan Klein's place on the ELSNET Board has now been taken over by **Geoffrey Sampson** from the University of Sussex. He was the third Board member to have been nominated via the new procedures, whereby ELSNET members are invited to propose candidates for vacancies.

Dr Sampson is Reader in Computer Science & Artificial Intelligence at the University of Sussex. His background is in NLP, but the integration of language and speech was already high on his agenda long before ELSNET came into existence: he organised the first conference in the UK explicitly dedicated to integrating NL and speech research (in 1987 at Leeds), and he was the founder and first director of the Leeds Centre for Computer Analysis of Language and Speech. All Sampson's NLP research since the early 1980s has made heavy use of corpora and statistical techniques, and he has a wide range of experience in the creation of language resources. The SUSANNE Corpus was entirely created under Sampson's direction, and he played a large role in



Geoffrey Sampson

the creation of the Lancaster Treebank. He is also familiar with issues in NLP evaluation. Sampson has worked as an industrial consultant on evaluation of NLP systems, for **British Telecom Research Labs** (now BT Labs) and for **Houghton Mifflin**.

In brief, Geoffrey has the ELSNET profile I believed only existed in dreams, and I am looking forward to collaborating with him! Geoffrey will also be Ewan's successor as convenor of the Information Dissemination Task Group, and will as such be responsible for our information and PR activities.

Welcome, Geoffrey!



A la recherche du temps perdu: members of one of the first ELSNET EB meetings. From left to right: (top) Gerrit Bloothoof, Joseph Mariani, Antonio Zampolli, Ewan Klein, Uli Heid, Niels Ole Bernsen; (bottom) Erik-Jan van der Linden, Isabel Trancoso, Anne De Roeck

e

Antonio Zampolli, University of Pisa

ELSNET's first beginnings precede Steven Krauwer's involvement in the network, but Antonio Zampolli was there. Here is his story.

Once upon a time — I don't remember exactly when — I was sitting in a bar in Brussels having a beer with an Italian representative on the ESPRIT Board (I was not practising the Montignac diet yet at the time). In the course of our conversation, this representative mentioned that the ESPRIT management would soon be launching a new type of initiative, called Networks of Excellence. He said it was not completely clear yet what the envisaged networks would look like, but they would be likely to provide new opportunities for collaboration between researchers and the exchange of doctoral students. I was interested, since the only real support from the Commission for international cooperation in the area of Language Engineering at the time was represented by EUROTRA; and I made a mental note of the information, before turning my attention to more immediate and more urgent things.

Shift to: Cambridge, 1991, the first review of ACQUILEX. Since ACQUILEX was the first European project I coordinated, and this was my first review, I was anxious and concerned about how it would be received by the reviewers, **Bente Maegaard** (whom I knew very well) and **Ewan Klein**, (whose name I was familiar with but whom I had never met). I had the idea that as a very formally-oriented linguist Ewan might not like the essentially data-driven approach of ACQUILEX too much, and I wanted to be very nice to him, make a good impression.

On the first day of the two-day review we went to the university restaurant to have lunch, and on our way there I found myself walking alongside **Simon Bensasson**, our project officer. Remembering the conversation in Brussels, I asked Bensasson about the Networks of Excellence Initiative. He explained what they had in mind and what a network would look like, and this made my mouth water... But he also pointed out that the game was nearly over, since the deadline for the pilot proposals was in three days' time. He added, however, that if we felt it was possible to prepare a proposal in such a brief time, the topic of integrating natural language and speech was, in his opinion, an important and promising one.

I thought things over and, before the start of the afternoon review session, I took aside Bente and Ewan, and suggested we should try to put together a proposal for a Network in Language and Speech. Ewan's immediate reaction was that it would not be possible to do this in the short time span available, and so I left it at that.



"Shall we split it 50:50?" Antonio Zampolli in conversation with Ewan Klein.

But at the social dinner that evening, Ewan didn't show up. I remember being concerned, asking myself whether his absence was due to the fact that he was not happy with the project. But half-way through dinner he called me to say he had stayed in to think about my suggestion. And the next day, after the conclusion of the review (a very positive one), he said that we should perhaps pursue the idea of a Network in Language and Speech after all. But he was concerned about the amount of time the management of such a network would absorb. My response to this was that we should allow in the budget for a half-time coordinator, supported by adequate infrastructure. But I didn't have the time to be that coordinator — it would have to be Ewan. He said he would think it over.

That evening Ewan told me that we should go ahead, and the next day we set to work. **Karen Sparck-Jones** gave us a room and two phones, we wrote a first draft of the proposal, and we established the criteria for institutes to join. Then we divided Europe into two parts, and phoned up the institutes that met our criteria to invite them to become managing nodes. All of them accepted, and by that evening both the structure of the network and the draft proposal were on the table. Over the next three days, Ewan and I were on the phone continuously. The rest of the story you know, of course.

In the realization of many initiatives I have been involved in, a series of coincidences seemed to play a decisive role; there's a story behind each of them. But one common determining factor was the fact that the right person was there at the right time. And so it was with ELSNET. It is clear that without Ewan, ELSNET would not have been possible at all. Not only because he played such a crucial role in writing the proposal, or because he had the authority to contact half of Europe; but also because he has an extraordinary capacity for conducting meetings, making everybody speak, gradually getting different positions to converge, and indicating the way towards priorities and targets. I particularly remember, with great admiration, how he coped with the original gap between the Speech and the Natural Language members in ELSNET, which he bridged with his diplomatic talents, his human qualities, his scientific competence and his breadth of

knowledge of both fields. Without these, ELSNET would not be what it is today — in fact it would simply not exist.

The fact that Ewan has (momentarily, I hope) retired from ELSNET is, to me, cause for nostalgia and regret: without Ewan ELSNET does not seem to be quite the same.

Antonio Zampolli is
Director of the Istituto di
Linguistica Computazionale
del CNR, Pisa.
Email: pisa@ilc.pi.cnr.it

ELSNET pages

Nov 1998

elsnet
.....

Smoke-filled rooms

Minutes of the ELSNET Executive Board meetings in June (Utrecht) and October (Madrid) summarised

Two ELSNET EB Meetings took place recently: the annual meeting of the ELSNET Extended Executive Board (consisting of both the Industrial Panel and the Executive Board members) was held in Utrecht, The Netherlands on June 21-22, 1998. And on October 18-19, the ELSNET Executive Board met in Madrid, where its members were warmly received by José Pardo and his staff at the department of Electronic Engineering of the Technical University of Madrid.

The following is a summary of the issues discussed at both meetings (which were, of course, entirely smoke-free).

New Executive Board member

Ewan Klein stepped down as member of the ELSNET Executive Board earlier this year, and ELSNET nodes were invited to propose candidates for EB membership. **Geoffrey Sampson** was appointed as new member on the basis of the new election procedure, which allows members to propose candidates, and which takes into account the profile and qualifications of the candidates, as well as the overall composition of the Board in terms of geographical distribution and thematic orientation. Sampson, from the University of Sussex (UK) was welcomed at his first EB meeting.

International Cooperation

A collaboration agreement was signed recently between the EU and the USA. This agreement offers an umbrella under which US and EU teams can now in principle collaborate under US and EU research programmes. ELSNET has plans for possible collaboration with the US in the field of cross-lingual text retrieval evaluation. Such collaboration could be extremely fruitful because it would combine shared interests with complementary experience and viewpoints.

ELSNET also has plans for cooperation with the Far East, and will fund an exploratory study to identify what data are already available and what data are needed, with the ultimate goal of a drawing up a survey of relevant actors in the Far East (including the Pacific Rim).

European projects

DISC

The DISC project (co-ordinated by The Maersk McKinney Moller Institute for Production Technology, Odense University) got a positive review. The project, which aims to establish best practice in the development and evaluation of Spoken Language Dialogue Systems and Components, runs for 18 months from 1 June 1997. It will be continued under DISC-II.

Contact person: Niels Ole Bernsen (nob@mip.ou.dk)

ELSE

ELSE is proposing a series of European-based evaluation competitions akin to the American MUC, TREC, PARSEVAL, and possibly others. So far six possible topics for separate competitions have been identified.

Contact person: Patrick Paroubek (pap@limsi.fr)

NAPLUS

ELSNET is involved in the NAPLUS project, coordinated by **Impetus Engineering Moschato** (Greece). The objective is to build a system for Arabic language understanding. Part of ELSNET's interest in this project is strategic: the hope is that it will lead to closer connections with the Mediterranean Language and Speech communities.

Contact person: Stelios Efstathiadis (impetus@compulink.gr)

Proposals for new actions


As always ELSNET members were invited to come up with new proposals. It was agreed at the meeting that ELSNET will support the following projects:

- **A Guide to the Language & Speech Domain in Russia:** this will be an entirely new (English) edition of the book "Les industries de la langue dans les pays de l'ex-USSR", which was issued and distributed in 1996 by the French Ministry of Science and Education. The book gives a survey of Russian teams and products in language and speech engineering. The project, coordinated by Vera Fluhr-Semenova, will be carried out by SCIPER (France).
- **Workshop on Prosody & Meaning.** This workshop, organised in collaboration with the MATE project, will be an additional event to the Vilem Mathesius Lecture Series at Charles University, Prague on 13-14 November 1998. For more information see <http://kwetal.ms.mff.cuni.cz/~gj/vmc/>
- Panel session at the **Senseval workshop** (2-4 September 1998), which aimed to integrate Senseval (see *ELSNNews* 7.3, p 8 for more on Senseval) into long term evaluation programs, particularly with respect to ELSE (cf. *ELSNNews* 7.3, p 13). For more information on the workshop, see <http://www.itri.brighton.ac.uk/events/senseval/>

New Nodes

The following sites were accepted as member nodes of ELSNET:

- University of Bergen (Norway). Contact person: **Koenraad De Smedt** (desmedt@hf.uib.no)
- University of Tbilisi (Georgia). Contact person: **George Chikoidze** (chiko@contsys.acnet.ge)



Recent ELSNET Executive Board Meetings



ELSNET EB members at the most recent meeting in Madrid

- Cap Gemini (The Netherlands). Contact person: **Jacqueline van Wees** (JWees@inetgate.capgemini.nl)
- Compuleer (The Netherlands). Contact person: **Marc Blasband** (cplr@worldonline.nl)
- ITRI, University of Brighton (UK). Contact person: **Adam Kilgarriff** (Adam.Kilgarriff@itri.brighton.ac.uk)
- SCIPER, Evry (France). Contact person: **Vera Fluhr-Semenova** (101376.156@compuserve.com)
- Alfa Informatica, Rijkuniversiteit Groningen (The Netherlands). Contact person: **Gertjan van Noord** (vannoord@let.rug.nl)

Multimodality & Multimedia

The EB unanimously agreed that multimodality & multimedia are very important topics for ELSNET. A range of ELSNET activities on these topics are planned, e.g. a workshop, multimodal resources and the 1999 ELSNET Summer School on Multimodality.

Training issues

The following issues were discussed:

- **EACL1999**

ELSNET and ELRA will organise a best paper award at EACL. EACL 1999 will be held from June 8-12 in Bergen, Norway. ELSNET members are invited to propose initiatives for ELSNET activities at EACL.

WEB: www.hit.uib.no/eac199/

- **Bullet course on Terminology**

The 2nd ELSNET Bulletcourse was organised by the HIT Centre at the University of Bergen from September 7-9. The topic was "The role of computer-based terminology management systems in translational information management". As the report elsewhere in this issue (p 3) shows, the event was very successful.

- **Eurospeech 1999**

ELSNET will organise a transliteration contest at Eurospeech 1999 (September 6-9 in Budapest, Hungary).

WEB: tel.ttt.bme.hu/Eurospeech99/

- **WEB-SLS**

After consultation with EACL and ESCA it was agreed that **Nikos Fakotakis** (University of Patras) will be the new editor of WEB-SLS, The European Student Journal of Language and speech.

WEB: web-sls.essex.ac.uk/web-sls/

2nd Annual Networks of Excellence Meeting

Since ELSNET's coordinator **Steven Krauwer** is president of the joint Networks of Excellence Coordinators Council, ELSNET will organise the 2nd Annual NoE Meeting in Vienna on 3 December 1998. More information about this meeting will be published in the next issue of *ELSNews*.

Next meeting

The next meeting of the ELSNET Executive Board will take place on March 7-8, 1999 in Greece. ELSNET members are invited to propose agenda points **before February 15, 1999**.

Nov 1998

elsnet
.....

Or Else ...

Yorick Wilks, *University of Sheffield*



We are delighted to welcome Yorick Wilks as a regular columnist for ELSNews over the coming year, and look forward to his critical, thought-provoking (and hopefully discussion-provoking!) pieces. Yorick kicks off in this issue with some reflections on EACL, the European Chapter of the Association for Computational Linguistics, and its role.

They asked me to do this and I just said yes; I didn't struggle, shuffle my feet, or need my arm twisted. Why, since, like you, dear Reader, I have many other things to do with my time? The truth is I rather admire ELSNET, even though I have had no previous direct connection to it, and now have no need to lick the hand that feeds me.

ELSNET, through ELSNews, seems to me to have established something very important in a dynamic research and development community: a relatively open forum, where researchers, users and finders can communicate with each other, inform and challenge each other, in an open manner. There is nothing like it, to my knowledge, in CL/NLP in the United States: the ACL (The Association for Computational Linguistics) does not play this role, though the AAAI (The American Association for Artificial Intelligence), by contrast, does very much provide it for artificial intelligence, interfacing with funders and representing users and researchers to them. Within NLP/CL in the US, ARP just announces its programmes and needs and that is pretty much that.

Without implying any criticism of anyone, I can tell you that I, like other people who were around at the time, had hoped that the EACL, the European Chapter of the ACL, would play this fuller role here.

It has clearly not done so, and it may be worth a moment to ask why, and to see if there are any general issues lurking nearby; doing that will also give me the chance to retail a great quote from the ACL-COLING 98 meeting in Montreal this summer.

I believe the reason the EACL was not able to develop the ELSNET role over the last years, in spite of all the other benefits it has brought to NLP/CL in Europe, is that its structure was, from the beginning, a colonialist one and it is only now beginning to escape from that and seek independence.

It was held back by a constitution that made it essentially a European branch office of the ACL in the US, and this was done deliberately at head office, and, as always, from the best possible motives; who could have known if the Europeans were fit to govern themselves? The Japanese politely declined when asked, ten years ago, if they wanted to set up an Asian branch chapter of the ACL.

Historians among you will remember the cry at US independence of "No taxation without representation", and it is a variant of this sentiment that has led to reaction against the present extraordinary situation, where European ACL members are almost half the total, but they are represented by a single delegate at the annual meeting. It was this relationship, this great historical reversal of fortune, with which we are all familiar in various ways, that has prevented EACL from fulfilling what might have been its local political role.

Those of you who were at the ACL general meeting in Montreal in August will already know that this may all be about to change. The ACL has realised the absurdity of all this and has proposed to the world membership a rational solution which has provoked very strong and varied reactions.

The proposed solution is a clean and consistent federal one: to set up a US chapter of the ACL, at the same level in the hierarchy as the existing European Chapter and any possible future Asian chapter. There would then be a core ACL, neutral between the chapters as it were, plus an

assumption that the main ACL conference would alternate between the regions/chapters.

These suggestions provoked some deep, atavistic, reactions, which can be epitomised by the American who argued plaintively, that the main ACL conference must continue to be held in the US, because of the high travel costs for Americans if it moved! [send me email if you don't get the joke at this point!] This point of view will remain a powerful one and may succeed in blocking any real change. After all, there will remain real problems beyond those of the symmetry of travel costs. If the main ACL conference rotates, will there then be local chapter conferences in the away-years? The split off of AAAI from the IJCAI conferences in artificial intelligence (whereby the AAAI was set up to hold a ISAI conference in years when IJCAI was not in North America) suggests there will be. If so, there will be a problem of whether papers can be submitted to all ACL conferences, as now, or only those run by one's own local chapter.

And one should not underestimate here the large number of cross-continental expatriates (e.g. the many US NLPers in Britain and elsewhere) who would not want to be cut off in a foreign, independent, chapter. Students of the break-up of the Roman Empire, or more recent post-colonial citizenship difficulties, will know the feeling.

All this remains unsettled, but its outcome will have a serious effect on the structure of the CL/NLP research community world wide, and to what degree it remains American-centred. Given the concerns and sensitivities, cultural and intellectual, of Europeans, this is an issue that matters very much to us. And it is why we should be glad and grateful that a determined group has set up a forum focussed on our concerns. Three cheers for ELSNET!

Yorick Wilks is Professor at the Department of Computer Science, University of Sheffield.
Email: yorick@dcs.sheffield.ac.uk
<http://www.dcs.shef.ac.uk/~yorick/>

e

Evaluation of NL projects and products

Bente Maegaard, Center for Sprogteknologi, Copenhagen.

... and another new feature: from this month onwards we'll be devoting some space in each issue to prominent or challenging issues in the area of Evaluation. Bente Maegaard wrote the first piece.

There has been a growing interest in evaluation methodologies in recent years, something which was very clearly demonstrated at the LREC conference in Granada (see *ELSN* 7.3 for more on LREC). Interestingly, evaluation problems have been tackled differently by the research communities working with speech and written language (NL) respectively. The speech community, working with very technical models, has been measuring performance using measurable criteria from the outset. Researchers in the NL area, on the other hand, were not concerned to the same extent with this type of evaluation some 20 years ago (with the notable exception of MT evaluation, which has been performed with varying degrees of success for 30 years!). But this has changed recently.

In Europe, interest in evaluation of NL products started to grow about 10 years ago, partly because researchers became aware of the ARPA/DARPA evaluation-guided research programmes in the US, which began running at that time. The main purpose of this type of programme, which promotes research through collaborative and competitive schemes, is to push research or technology development in the area being evaluated (e.g. text retrieval), instead of research in evaluation itself. Another (and probably more important) reason for the recent upsurge in interest in NL evaluation in Europe lies in the fact that language technology is now so advanced, and language technology products are becoming so widespread, that there are now consumers who want to know which system to buy.

The ARPA/DARPA type of evaluation has been much discussed, as it has its merits and drawbacks. In Europe similar efforts exist, for example the French AUPELF/UREF programme; but until recently most efforts in the area of evaluation in Europe focussed on the methodology itself, most notably through the EAGLES project. The EAGLES project had two main goals: to develop a general evaluation methodology which can be used for all types of NL evaluation; and to give special consideration to evaluation with respect to user needs and usability. Standardisation is hard to achieve, as it is always difficult to reach consensus, but it seems that the EAGLES effort has been successful; and at LREC it became clear that there is a lot of interest from American researchers and funding agencies too in using and giving feedback on the EAGLES methodology. So it seems that a consensus can be built, not only in Europe but also across the continents, on how to think about evaluation, what the important questions are, and how reliable evaluations can be set up that answer these questions.

The purpose of an evaluation effort will differ slightly depending on who needs the evaluation: a researcher will want to know if his or her technology has improved, or if it meets the goals set for it, and a funding agency may have similar goals; but an end user will have a different perspective. An end user cannot be satisfied with a technology evaluation, but will need to take



Bente Maegaard

into account the task for which the technology is going to be used. So a user-centred evaluation will start from the actual *use* of the system for a certain task by a certain user — or rather a certain class of users.

Performing an evaluation basically involves three tasks, namely

1. definition of the elements of the evaluation (system, user, measures and methods for measuring);
2. collection of test materials, and
3. running the testing and compiling the report.

Task 1 is the most difficult one. For simple NL products it may be reasonably straightforward to describe the intended functionality, usability criteria and so on, and to break them down into sub-criteria until a criterion is reached which can be measured in a well-defined way. But for more complex systems it can be quite complicated, as is clear from the MT evaluations carried out over the years. Part 2 is time-consuming, but as language resources are being built they can hopefully be reused (directly, or after modification) for several purposes, including evaluation. Part 3 should be easy, in particular if automated tools have been built.

Future developments will hopefully be collaborative, in the sense that they strengthen the emerging consensus on evaluation methodology and allow for test data and evaluation tools to be shared. After LREC, I think that there are good reasons to believe this can and will happen.

Bente Maegaard (bente@cst.ku.dk) is director of the Center for Sprogteknologi, Copenhagen.

Nov 1998

elsnet
.....

GATE to infrastructure

Hamish Cunningham, University of Sheffield

If you're a professional mathematician, you probably regard a tool like SPSS or Mathematica as necessary infrastructure for your work. But if you're a computational linguist or a language engineer, chances are that large parts of your work have no such infrastructural support, and attempts to provide one over the last decade have not been very successful. GATE, developed over the last three years at the University of Sheffield, aims to fill this gap.

What does infrastructure mean for Natural Language Processing (NLP)? What sorts of tasks should be delegated to a general tool, and which should be left to individual projects? GATE does three things:

- it manages textual data storage and exchange;
- it allows easy visualisation of textual data structures; and
- it provides plug-in modularity of text processing components.

Based on the collective experiences of a sizable user base across the EU and elsewhere, the system can claim to be the Mathematica of NLP for certain sections of the field. Given further development, we hope that it can take this role for a wide variety of tasks.

There are many useful tools around for performing specific tasks such as developing feature structure grammars for evaluation under unification, or collecting statistical measures across corpora. To varying extents, these entail the adoption of particular theories, which limits their take-up. The only common factor of NLP systems, alas, seems to be that they all create information about text. Developers of such systems create modules and data resources that handle text, and they store this data, exchange it between various modules, compare results of test runs, and generally spend inordinate amounts of time pouring over samples of it when they really should be enjoying a slurp of something relaxing instead.

The types of data structure typically involved are large and complex, and without good environments to manage succinct viewing of the data we work below our potential. At this stage in the progress of our field, no one should really have to write a tree viewing program for the output of a syntax analyser, for example, or even have to do significant work to get an existing viewing tool to process their data.

In addition, many common language processing tasks have been solved to an acceptable degree by previous work and should be reused. Instead of writing a new part of speech tagger, or sentence splitter, or list of common nominal compounds, we should have available a store of reusable tools and data that can be plugged into our new systems with minimal effort. Such reuse is much less common than it should be, often because of installation and integration problems that have to be solved afresh in each case.

An architecture is thus a macro-level organisational pattern for the components and data resources that make up a language processing system; development environments add graphical tools to access the services provided by the architecture. The

GATE architecture does this using three subsystems:

- GDM, the GATE Document Manager;
- GGI, the GATE Graphical Interface;
- CREOLE, a Collection of REusable Objects for Language Engineering.

GDM manages the information about texts produced and consumed by NLP processes; GGI provides visual access to this data and manages control flow; CREOLE is the set of resources so far integrated. A developer working with GATE begins with some subset of CREOLE that do some basic tasks, perhaps tokenisation, sentence and paragraph identification and part-of-speech tagging. They then add modules for their specific tasks. They use a single API for accessing the data and for storing their data back into the central database. With a few lines of configuration information they allow the system to display their data in friendly graphical form, including tree diagrams where appropriate. The system takes care of data storage and module loading, and can be used to deliver embeddable subsystems by stripping the graphical interface. It supports modules in any language including Prolog, Lisp, Perl, Java, C++ and Tcl.

The GATE system is used in several Framework IV projects and installed at over 200 sites world-wide. At the LREC workshop in Granada earlier this year we presented plans to incorporate resource distribution, examination, use and evaluation into GATE. It is being adopted by the US TIPSTER program as part of its **Architecture Capabilities Platform**, and is available free of charge on licence from Sheffield.

FOR INFORMATION

More details on GATE may be found at

dcs.shef.ac.uk/research/groups/nlp/gate/

and a growing set of available resources at

<ftp://ftp.dcs.shef.ac.uk/share/nlp/CREOLE>


Plans for the future are at

dcs.shef.ac.uk/~hamish/GATE2.html

Hamish Cunningham is a Research Fellow in Computer Science at the University of Sheffield.

Email: hamish@dcs.shef.ac.uk

Web: dcs.shef.ac.uk/~hamish/



Niels Ole Bernsen, Hans Dybkjær and Laila Dybkjær: *Designing Interactive Speech Systems*. Springer, 1998

Review by **Jeremy Peckham**.

Interactive speech systems are beginning to impinge more and more on many aspects of our lives, whether it be making a directory enquiry, checking a bank account balance or finding a train time. *Designing Interactive Speech Systems* is the first coherent and extended text which I have come across which seeks to address the vexed issue of system design and evaluation.

The creation of good spoken dialogue systems is at present as much art as science. Niels Ole Bernsen, Hans Dybkjær and Laila Dybkjær have made a significant contribution to the science by seeking to provide some much-needed structure to the process. Covering the issue of assessing what the speech modality is good for, speech interaction theory, through to design principles, wizard of oz simulation and system evaluation, the authors provide insights from their experiences in the *Danish Dialogue Project*. At the heart of the text are useful guidelines for designing co-operative dialogues, a feature the authors consider to be key to successful interaction. These are compared with Grice's maxims, and also tested against user experience with the flight reservation domain of the Danish Dialogue Project. Given the state of the art in speech systems, the authors contention that dialogue repair and 'meta-communication' skills form an important part of co-operative systems is well made.

Two areas that I would like to have seen developed further are requirements specification and 'adequacy evaluation'. Although these areas are touched on, one of the biggest commercial and research challenges is finding the matches between business applications, user expectations and current technology limitations. As the authors point out, much of this activity requires the skills and knowledge of expert developers.

Whilst there is still much research to be done in spoken dialogue systems, this book should be required reading by all those attempting interactive speech system design, whether researchers or commercial implementers.

Jeremy Peckham (JBPeckham@aol.com) is Managing Director of Strategis Consulting Ltd

Gerrit Bloothoof et al. (eds): *The Landscape of Future Education in Speech Communication Sciences. Vol. 2: Proposals*. Utrecht Institute of Linguistics OTS, 1998.

Review by **Klaus Fellbaum**.

Since 1996, the SOCRATES Thematic Network *Speech Communication Sciences* has been developing visions of new directions for academic education in the Speech Sciences. The Network, which brings together about 100 university departments in Europe, recently published a second book in its series *The Landscape of Future Education*.

While the first book (published in 1997) analysed the current state of education in speech communication sciences in Eu-

rope, the new book offers a number of proposals for academic studies in phonetics, spoken language engineering (SLE) and — a new element — speech and language therapy. It also has a section on computer-aided learning and the Internet.

The section on spoken language engineering discusses the results of a questionnaire amongst SLE professionals on the relation between education and job requirements. One significant result was that in many areas job training and/or personal effort were rated as significantly more important than university education. This suggests that the number of university courses in speech and language engineering need to be increased, and their content improved.

The book goes on to present an outline for curricula in SLE. The work of a SL engineer not only concerns technological areas like coding, recognition and synthesis, but also involves signal processing and programming skills, as well as disciplines like phonetics and linguistics. To meet the different requirements, the SLE curriculum is divided into seven levels. Lectures and other teaching activities (from the teaching staff) are widely replaced by independent but supervised research-oriented work combined with seminar presentations.

Other issues discussed are the European Masters in Language and Speech currently being developed (presented in more detail on pp. 1-2 of this issue), and the integration of speech communication sciences in speech and language therapy curricula. The latter proves to be an exciting issue, since up till recently speech and language therapy was a more or less 'closed' area of clinical therapy experts. Descriptions in the book show that contacts between representatives from speech communication engineering, linguistics and phonetics has brought about a synergy valuable to all.

Finally, the book contains a lot of advice on computer-aided learning and the use of the Internet in education. It discusses a wide range of applications, from simple course pages to very complex Web-based interactive tutorials. It also has a very useful inventory of reading material and Internet supports.

It may be worth pointing out that the material presented in this book is, as the title suggests, intended as a proposal. The authors emphatically invite, and hope to receive, many comments (including critical ones) — something which would demonstrate again that the field of speech and language is very much alive and full of activity.

Klaus Fellbaum (fellbaum@naxos.kt.tu-cottbus.de), Professor

Lehrstuhl Kommunikationstechnik, Cottbus Technical University Cottbus

FOR INFORMATION

For more information on the Thematic Network **Speech Communication Sciences**, see www.tn-speech.essex.ac.uk/tn-speech/ or contact **Gerrit Bloothoof** (Gerrit.Bloothoof@let.uu.nl). The book can be obtained for free from the same address.

Book reviews

Nov 1998

elsnet
.....

First announcement

7th European Summer School on Language and Speech Communication Multimodality in Language and Speech Systems (MiLaSS)

Stockholm, Sweden, 12-23 July 1999

The 1999 ELSNET Summer School will be hosted by the Department of Speech Music and Hearing at Kungliga Tekniska Högskolan (KTH) in Stockholm. The event will be sponsored by ELSNET, TMR and KTH, and receive further support from ESCA and EAEL.

For details of the programme, which will consist of plenary sessions, parallel courses and workshops, please check the Summer School homepage:

www.speech.kth.se/milass/

Courses at the school are aimed at doctoral students, and the number of students will be limited to 60, so early registration is recommended. Pre-registration will start late autumn 1998. A registration form will soon be published on the Summer School homepage.

FOR INFORMATION

For more information on the school, contact:
MiLaSS

Dept of Speech, Music and Hearing, KTH
S-100 44 Stockholm

Sweden

Tel: +46 8 790 7879

Fax: +46 8 790 7854

email: MiLaSS@speech.kth.se

Web: www.speech.kth.se/milass/

Both ESCA and TMR grants will be available. For information on TMR grants, see

www.speech.kth.se/milass/

For information on ESCA grants, see

www.esca-speech.org/grants.html

*Call for papers***MATISSE**

ESCA/SOCRATES workshop on

Method and Tool Innovations for Speech Science Education

University College London, April 16-17, 1999

This joint ESCA/SOCRATES workshop aims to bring together educators and developers of educational tools and materials in the Speech Science area. The focus of the workshop will be on innovative approaches and tools for teaching speech sciences (phonetics, spoken language engineering, speech and language therapy). Contributions are invited which describe innovations in the field of:

- teaching methods
- teaching resources
- computer-aided learning
- evaluation
- curricula

The session on curricula will involve a discussion of proposals put forward by the SOCRATES Thematic Network, to be published in September 1998.

Technological demonstrations will be particularly welcome.

FOR INFORMATION

Further information about the MATISSE workshop is available at the MATISSE web site:

www.phon.ucl.ac.uk/home/matisse/first

The deadline for abstract submission is 15 December 1998.

The workshop will be preceded by the first meeting of the *Phonetics Teaching and Learning Conference*, which will take place at University College London on 14th and 15th April. This conference is organised by the SIPhTrA project (System for Interactive Training and Assessment). For details of this please contact **John Maidment** (johnm@phonetics.ucl.ac.uk) or see the pTLC99 home page at

www.phon.ucl.ac.uk/home/johnm/ptlc.htm

Speech Recognition Update (SRU)

Special SRU subscription discount for ELSNET Members

Speech Recognition Update (edited by William S. Meisel) is a monthly newsletter with news and analysis of speech recognition markets, products and technology.

ELSNET members will receive a 20% discount on the first year of their subscription to the newsletter.

Yes, I am an ELSNET member and I wish to subscribe to **Speech Recognition Update** for one year (12 issues) at the following reduced rate:

US \$ 384 International rate (incl. Canada) (full rate US \$ 480)

US \$420 International Library rate (full rate US \$ 525)

Name _____

Organisation _____

Address _____

City/Postal Code _____

Country _____

Tel _____

Fax _____

Check enclosed (in US\$ on a US bank)

Bill me

Bill my Visa/MasterCard/AmericanExpress

Expiry Date _____

Signature _____

Please mail or fax this form to:

*TMA Associates,
P.O. Box 570730,
Tarzana, CA 91357-0730
USA.
Tel: (818) 708-0962.
Fax: (818) 345-2980*

Future events

Nov 30-Dec 2, 1998: *IST 98* Conference and Workshops in Vienna, Vienna, Austria. Further info: Roberto.Cencioni@cec.be, <http://www.cordis.lu/ist98/>

Nov 30-Dec 4, 1998: *ICSLP 98*, Sydney, Australia. Further info: icslp98@tourhosts.com.au, <http://cslab.anu.edu.au/icslp98>

Dec 7-8, 1998: *TWLT14*: Language technology in multimedia information retrieval, Twente, Netherlands. Further info: netter@dfki.de, <http://www.seti.cs.utwente.nl/Parlevink/Conferences/twlt14.html>

Dec 8, 1998: *Knowledge-Based Language Engineering*. Commonwealth Institute, London. Further info: unicom@unicom.co.uk

Dec 14-16, 1998: *LACL 98* (Logical Aspects Of Computational Linguistics), Grenoble, France. Further info: lacl98@let.ruu.nl, <http://www.loria.fr/~retore/LesLACL.html>

Dec 17-20, 1998: *LOFT3* conference, Torino, Italy. Further info: icer@inrete.it, <http://pages.inrete.it/icer>

Jan 4-27, 1999: *1999 Winter School in Tver*, Tver, Russian Federation. Further info: Email: inforuss@postman.ru

Jan 11-22, 1999: *Winterschool 1999*, Amsterdam, Netherlands. Further info: lot@let.uu.nl, <http://www.lot.let.uu.nl/ws99.html>

Jan 13-15, 1999: *IWCS-3*, Tilburg, Netherlands. Further info: Computational.Semantics@kub.nl, <http://cwis.kub.nl/~fdl/research/ti/Docs/IWCS/iwcs.htm>

Feb 17-19, 1999: *International Conference on Computational Intelligence for Modelling*, Vienna, Austria. Further info: cimca99@fcit.monash.edu.au, <http://www-gscit.fcit.monash.edu.au/conferences/cimca99>

Mar 12-14, 1999: *2nd International Conference on Voice Physiology and Biomechanics*. Berlin, Germany. Further info: cranen@let.kun.nl, <http://www.nwu.edu/csd/ICVPB/ICVPB.html>

Apr 14-15, 1999: *Phonetics Teaching and Learning Conference*, University College London, United Kingdom. Further info: val@phonetics.ucl.ac.uk, <http://www.phon.ucl.ac.uk/home/johnm/ptlc.htm>

June 8-12, 1999: *EACL 1999*, Bergen. Further info: <http://www.hit.uib.no/eacl99/>

Sep 5-10, 1999: *Eurospeech 1999*, Budapest. Further info: <http://tel.ttt.bme.hu/Eurospeech99/>

Special offer

Future events

Nov 1998



ELSNET Secretariat

Steven Krauwer
Coordinator

Mariken Broekhoven
Assistant Coordinator
Utrecht University (NL)

Task Group Convenors

Training & Mobility
Gerrit Bloothoof, Utrecht University (NL)

Info Dissemination
Geoffrey Sampson
University of Sussex (UK)

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

Research
Niels Ole Bernsen
Odense University and Joseph Mariani
LIMSI-CNRS

Industrial Panel

Harri Arnola, Kielikone (SF)
Roberto Billi, CSELT (I)
Michael Carey, Ensigma (UK)
Jean-Pierre Chanod, Rank Xerox Research Centre (F)
Harald Höge, Siemens AG (D)
Bernard Normier, ERLI (F)
Brian Oakley (chair, UK)

ELSNET Participants

Academic Sites

A OFAI/Univ. Vienna/Vienna Univ. of Technology
B University of Antwerp
B University of Leuven
BU Bulgarian Acad. of Sciences, Sofia
BY Belarussian Academy of Sciences, Minsk
CH IDSIA, Lugano
CH ISSCO, Geneva
CZ Charles University, Prague
D Univ. des Saarlandes, Saarbrücken
D DFKI, Saarbrücken
D IAI, Saarbrücken
D Univ. Hamburg
D Univ. Kiel
D Univ. of Stuttgart
D Ruhr-Univ. Bochum
D Univ. Erlangen
DK Ctr for Sprogteknologie, Copenhagen
DK Ctr for PersonKommunikation (CPK), Aalborg
DK Odense University
E Universidad de Granada
E Univ. Politecnica de Catalonia/Univ. Autonoma de Barcelona
E Univ. Politecnica de Madrid
E Univ. Politecnica de Valencia
F LIMSI-CNRS, Orsay
F IRIT, Toulouse
F Inst. de la Comm. Parlée, Grenoble
F IRISA, Rennes
F Laboratoire Parole et Langage-CNRS, Aix-en-Provence
F CRIN, Nancy
G University of Tbilisi
GR ILSP/NCSR "Demokritos", Athens
GR Wire Communications Lab., Patras
H Hungarian Acad. of Sciences, Budapest
H Technical University, Budapest
I Ist. di Linguistica Computazionale, Pisa
I IRST, Trento
I Fondazione Ugo Bordoni, Rome

IRL University College Dublin
IRL University of Dublin
LT Institute of Mathematics and Informatics, Vilnius
N University of Bergen
N University of Trondheim
NL Stichting Spraaktechnologie, Utrecht
NL Inst. for Perception Research, Eindhoven
NL Leyden Univ.
NL Catholic Univ. of Nijmegen
NL Rijksuniversiteit Groningen
NL TNO Human Factors Research Institute
NL Univ. of Amsterdam
NL Univ. of Tilburg
NL Univ. of Twente
NL Utrecht University (coordinator)
P INESC/ILTEC/Univ. Nova de Lisboa
PL Polish Academy of Sciences, Warsaw
RO Research Inst. for Informatics, Bucharest
RU Russian Academy of Sciences, Moscow
S KTH, Stockholm
S Univ. of Linköping
UK ITRI, University of Brighton
UK Defence Research Agency, Malvern
UK UMIST, Univ. of Manchester
UK Univ. of Cambridge
UK Univ. College London/School of Oriental and African Studies (SOAS)
UK University of Edinburgh
UK Univ. of Essex
UK Univ. of Dundee
UK Univ. of Leeds
UK Univ. of Sheffield
UK Univ. of Sunderland
UK Univ. of Sussex
UK Univ. of Ulster
UK Univ. of York

D Langenscheidt
D Novotech GmbH
D pc-plus Computing
D Philips Research Laboratories
D Siemens AG
D Verlag Moritz Diesterweg
DK Tele Denmark
E Telefonica I&D
F ACSYS
F Aerospatiale
F GSI-ERLI
F LINGA s.a.r.l.
F MemoData
F Rank Xerox Research Center
F SCIPER
F Systran SA
F TGID
F VECSYS Speech Processing
GR Knowledge A.E.
H Morphologic
I CSELT
I Database Informatica
I Sogei (IRI-FINISIEL Group)
I Syntax Sistemi Software
I Tecnopolis CSATA Novus Ortus
I Olivetti Ricerca SpA
NL Cap Gemini
NL Compuleer
NL KPN Research Laboratories
NL Polydoc N.V.
RU Analit, Ltd.
RU Russicon Company
S Telia Promotor (Call Centre Division)
FIN Nokia Research Center
FIN Kielikone Ltd
UK ALPNET UK, Ltd
UK BICC plc
UK British Telecommunications
UK Cambridge Algorithmica Ltd.
UK Canon Research Centre Europe Ltd.
UK Ensigma Ltd.
UK Hewlett-Packard Labs
UK Logica Cambridge Ltd.
UK Sharp Laboratories
UK SRI International
UK Vocalis Ltd.

Industrial Sites

B Lernout & Hauspie Speech Products
D aspect GmbH
D Daimler-Benz AG
D Electronic Publishing Partners GmbH
D Grundig Professional Electronics GmbH
D IBM Deutschland

What is ELSNET?

ELSNET, the European Network in Language and Speech, was established in 1991 with funding from ESPRIT Basic Research. There were 25 founding members of the network. Currently, there are more than 60 universities and research institutes, and more than 45 companies participating.

The long-term technological goal which unites the members of ELSNET is to build integrated multilingual NL and speech systems with unrestricted coverage of both spoken and written language. Building multilingual NL and speech systems requires a massive joint effort by two pairs of communities: on the one hand, the natural language and speech communities, and on the other, academia and industry. Both pairs of communities are traditionally separated by wide gaps.

It is ELSNET's objective to provide a platform which bridges both gaps, and to ensure that all parties are provided with optimal conditions for fruitful collaboration. To achieve this, ELSNET has established an infrastructure for sharing knowledge, resources, problems, and solutions by offering (information) services and facilities, and by organising events which serve academia and industry in both the language and speech communities.

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 persons who have an interest in ELSNET's activities. This mailing list may be used to announce activities, post job openings, or discuss issues which are relevant to people in the European natural language and speech communities. To request additions/deletions/changes of address in the mailing list, please send mail to elsnet@let.ruu.nl.

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@let.uu.nl
Web: www.elsnet.org