

The ICSI GAZETTE

volume five | issue two | march 2007

featured research: spoken language processing

in this issue:

page 2
As I See It
by Nelson Morgan, Director

page 3
News Briefs

page 4
Featured Alumni

page 9
Visiting Scholars

page 10
Publications

ICSI's Speech Group has grown significantly since its early days as the Realization Group. The group continues to stay true to its roots as an innovator in the field of automatic speech recognition (ASR) with a unique focus on fostering worldwide collaborations. At the same time, the research scope of the group is expanding and now includes speaker recognition, speech processing of multiparty meetings, dialog systems and information distillation, and porting of ASR systems to work in multiple languages. A recent project of the Speech Group also involved making speech technology accessible and useful to illiterate people in third world villages (see September 2006 Gazette).

As the group nears the end of two decades of speech research, the success of the collaborative aspect of speech research at ICSI is striking. Hervé Bourlard, one of the first visitors to the speech group, now heads IDIAP, a close collaborator with ICSI on speech processing for meetings, and also serves on ICSI's Board of Trustees. He is just one Speech alum who continues to contribute - alumni from the group have a way of finding themselves back at ICSI for research visits. Current staff scientists Nikki Mirghafori, Andreas Stolcke, and Chuck Wooters are all group alumni. Dan Ellis, one of the featured alumni in this issue, continues to work with ICSI on meeting research, although he is an associate professor at Columbia University in New York. He is one of many group alumni who have gone on to create collaborations between ICSI and their new university or lab. Our other featured alum, Eric Fosler-Lussier, continues to develop speech processing techniques he first worked on while at ICSI. In addition, thriving visitor programs continue to provide new insights and perspectives of young, up-and-coming scientists who visit the Speech Group from overseas.

Global Autonomous Language Exploitation: GALE has the goal of improving translation and interpretation of input from multiple languages in spoken and written form.

Our featured project for this issue, the Global Autonomous Language Exploitation (GALE) project embodies the spirit of collaboration at ICSI as well as the quality of research, and is challenging the speech staff to pursue exciting new directions in spoken language processing.

GALE's ambitious goals require significant improvement of speech recognition, diarization, sentence segmentation, machine translation, and information distillation in English, Mandarin Chinese, and Arabic.

ICSI's participation in GALE is part of a large collaborative effort called Nightingale led by SRI and including researchers from 15 sites, including IDIAP, University of Washington, and Columbia University. The SRI team is one of three working on the GALE effort; the other two are led by BBN and IBM. We focus here on ICSI's contributions to Nightingale.

AUTOMATIC SPEECH RECOGNITION (ASR)

ICSI researchers applied their proven ASR methods used for American English to Mandarin and Arabic with very promising results. These methods

Continued on Page 11



as i see it by Nelson Morgan, Director

The focus in this issue on speech research reminds me of my early days at ICSI, almost 2 decades ago. The Institute was also just starting out, and I had been asked to help put together a “Realization” group that would design and build massively parallel systems. This sounded like fun, but I also wanted to have an application area to focus on so that our work would not be an abstract exercise. At about that time Hervé Bourlard came to ICSI as a visitor, and “infected” us with his enthusiasm for speech recognition research. I had worked in speech processing earlier in that decade, and after several years of working with brain scientists to understand a little bit about the neurophysiological correlates of cognition, I was ready to return to speech, at least for a while.

Hervé had been working at Philips in Brussels, and had realized that one could estimate probabilities with properly trained neural networks. Both his theoretical work and his intuition told him that these probabilities could be used to improve speech recognition if incorporated with hidden Markov models. This sounded intriguing to me, so I set to work with him on the related experimental research. Our first result with a problem of scale was very motivating – we got 140% error! To the uninitiated this might sound impossible, but it turns out that in speech recognition we count as errors not only the words that are wrong or missing, but also the extra words that are inserted in the output. Anyway, we figured we had nowhere to go but up ... (or down, if you’re counting errors). And indeed we did. Over the months that followed, we gradually discovered, step by step, the things we needed to do for good performance. The next year we were joined by Chuck Wooters, probably the only student to ever get a degree at Berkeley in the interdisciplinary cross-department topic of “speech recognition.” With Chuck’s help, and in a collaborative effort with Mike Cohen and Horacio Franco (both then of SRI), we ultimately ended up with a very good system and plenty of new ideas.

Another milestone from those early days came from our collaboration with Hynek Hermansky. At that time Hynek was working for US West (later Qwest). Unlike Hervé, who was primarily concerned with the statistical models, Hynek was sharply focused on the feature extraction process. Previously he had invented a technique called “Perceptual Linear Prediction” or PLP, which is now used by many systems. He had been motivated by the goal of making systems more independent of variation in the signal due to different speakers, but when

we started working together he was interested in making recognizers less sensitive to other kinds of variability. Our interest in this topic was piqued by a small workshop we organized (the SPeech recOgnition frOnt eNd workshop, or SPOONS), for which we invited a number of people who had designed innovative models for speech processing, e.g., Les Atlas, Jordan Cohen, Ron Cole, Malcolm Slaney, Dick Lyon, Dirk Pueschel, and Shihab Shamma. The discussions were great (Mike O’Malley, who had done important work in speech synthesis, asked “Why a 10 millisecond analysis step? Because we have 10 fingers?”), but one of the comments particularly struck Hynek and me. Jordan Cohen, who had designed a biologically-inspired speech recognition front end for IBM, asked: “We could play speech through a filter approximating the inverse of a steady state vowel spectrum (such as ‘e’) and the speech is still intelligible, including the vowels which turn into a white spectrum signal. Which hearing model can account for that?” Given this “inverse-e” challenge, Hynek and I later came up with what we called RelATive SpecTRal Analysis, or RASTA, an approach that ultimately was adopted by Qualcomm and ended up in many millions of cell phones as the front end for speech recognition.

In the following years, I was fortunate to work with both Hervé and Hynek as they changed jobs. As our group grew, adding students, postdocs, and the continuing flow of talented visitors, we all worked closely with these two early contributors and the teams that worked with them, for instance co-developing new approaches to the incorporation of multiple feature streams in speech recognition. As of this writing both of them are at IDIAP in Switzerland (in many ways our sister institution) and our joint work continues.

Since those early days, we have graduated 16 PhDs from the Speech (née Realization) group, some of who are now teaching a new generation of students; and others have come back to ICSI as research staff. We now have a new generation of the next 10 students who are working on a much more diverse set of problems, as we have expanded from speech recognition to speaker recognition, sentence segmentation, “diarization” (who spoke when), and a number of aspects of speech understanding. I hope that this issue of our Gazette, which will include a focus on the DARPA-sponsored GALE project, will give our readers some insight as to the directions of the current group.

news briefs

WINTER 2006-2007

The **SCALE** vector-thread microprocessor prototype, developed by a team from MIT including recent ICSI visitors Professor **KRSTE ASANOVIC** and his students **RONNY KRASHINSKY** and **CHRIS BATTEN**, was completed in February. The device was a winner in the 2007 ISSCC/DAC (International Solid State Circuits Conference/Design Automation Conference) student design contest, and will be presented at the upcoming DAC conference this summer.

Congratulations to **ERAN HALPERIN** of the Algorithms Group and his wife Leticia. Their son Yanai was born on January 29th. Yanai weighed 7.3 pounds and was 20.5 inches long at birth.



Yanai Halperin

ICSI's **PAUL KAY**, of the AI Group, was interviewed on NPR's "All Things Considered" on Saturday, January 27th. An audio file of the story, "Developing a Vocabulary of Color", is available from NPR's web site. <http://www.npr.org/templates/calendar/index.php?prgId=2>

Congratulations to **PAVLIN RADOSLAVOV** of the Networking Group and his wife Kazuyo on the birth of their daughter Maria on January 24th. Maria was 6 pounds, 11 ounces at birth and 19 1/2 inches long.

In honor of **NORBERT SZYPERSKI** on his 75th birthday, ICSI has created a fund in his name. Szyperski contributed to the

founding of ICSI and his personal efforts over the years on behalf of the Institute are remarkable. It was his vision and his sustained endeavor that made the success of the Institute possible. ICSI is now accepting contributions to the fund, which may qualify as tax-deductible within the U.S.

ICSI research made the news at www.economist.com on January 18th. An article in the science and technology section titled "How grue is your valley?" discusses recent research by ICSI's **PAUL KAY**, ICSI alum Terry Regier (now at University of Chicago) and their colleagues, on the relationship between color and language in the human brain.

VERN PAXSON of the Networking Group is now a Fellow of the ACM (Association for Computing Machinery). A January 8, 2007 press release from the ACM named Paxson as one of 41 fellows for 2006. He was recognized for his contributions to Internet measurement and intrusion detection.

Congratulations to **MARK ALLMAN** of the Networking Group and his wife Meredith on the birth of their daughter, Sarah. Sarah was born on January 2nd. She weighed 7 pounds 12 ounces and was 20 inches long at birth. Sarah is the Allmans' second daughter.

FALL 2006

Congratulations to **LIZ SHRIBERG** and **ANDREAS STOLCKE** of the Speech Group on the birth of their son, Peter Liev Stolcke, on November 8th. Peter was 6 pounds 13 ounces at birth and 19.5 inches long. Peter's four older brothers are excited to finally complete their basketball team.

Professor **RICHARD KARP**, head of the ICSI Algorithms Group, was quoted in an October 31 New York Times article by Steve Lohr. Karp discussed the relationship between computer science and biology, pointing out that algorithms are good at describing dynamic biological processes, like protein production. Karp's presentation at the recent "2016" symposium clearly made an impression on Lohr, who writes, "[Karp] presented a fundamental explanation for why computing has had such a major impact on other sciences, and Dr. Karp himself personifies the trend. His research has moved beyond computer science to microbiology in recent years".



Professor Rickard Karp

Mark Ashida discussed **XORP**, ICSI's eXtensible Open-source Router Project, in a podcast posted on the Port 25: Communications from the Open Source Lab @ Microsoft web site on Thursday, October 26.

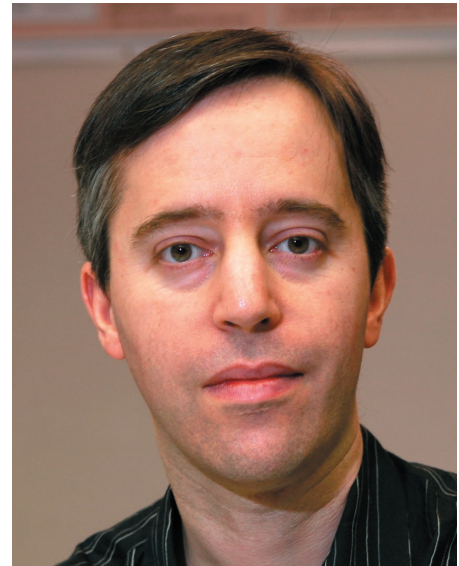
Professor **PAUL KAY** was interviewed for an October 16 Nature magazine article reporting on a recent study by scientists at Ohio State University on universals of color naming. The current study uses data from the World Color Survey to make a strong statistical case in support of ideas long held by Kay and his colleagues on color naming.

Continued on Page 6

featured alum: dan ellis

Dan Ellis is an associate professor at Columbia University in New York City. His research focus is signal processing and machine learning techniques, which he applies to speech recognition and music analysis. He received his PhD in electrical engineering at MIT. Upon completion of his degree, he worked at ICSI as a postdoctoral researcher from 1996-1999, then as a staff research scientist before being offered his current position at Columbia in 2000. Ellis continues to maintain an active affiliation with ICSI. Since leaving ICSI, he has worked with ICSI's Speech Group on the EARS project as well as on speech processing for meetings, and he is involved in this issue's featured project, GALE.

At ICSI, Ellis researched signal processing techniques, with a particular focus on using signal processing for automatic speech recognition (ASR). Through his ASR work he had the opportunity to collaborate with many different people, including Professors Nelson Morgan and Hervé Bourlard. He says that collaboration was one of the most valuable things he learned at ICSI. Its success at enhancing the quality of research there. This instilled



ICSI Alum Dan Ellis

“At ICSI I learned how to collaborate, and how to make collaborative projects successful, which has had a big influence on my career.”

--Dan Ellis, Columbia University

in him a strong belief in the value of collaborative research, and his continued cooperation with ICSI and other institutions has had a strong influence on his career thus far. He feels that understanding how to make collaborative projects work is a valuable skill as an academic.

Although speech technology continues to be a big part of Ellis's research, he recently began some work using some of the same signal processing and machine learning techniques to identify features in music. In one current project, he trains an algorithm to recognize chords in

music. Chords are a basic unit for music, analogous to the words in speech. Although there is much more work to do, such as increasing the amount of training data and defining which features are important, the results are promising so far. He hopes to develop applications for music signal processing, such as information retrieval, creating customized playlists for users, and determining new music that a user might like based on its similarity to the user's known preferences. Ellis teaches courses on music signal processing and machine learning at Columbia, in addition to general signal processing and speech and audio signal processing.

Combining his affinity for collaboration as well as teaching, Ellis has been working with two of ICSI's graduate student researchers, Arlo Faria and Kofi Boakye. Faria contacted Ellis while updating the Quicknet software originally developed at ICSI by Ellis and Dave Johnson. It is being used as front end for ASR in the GALE project, and Faria is adapting it to work more quickly on ICSI's current server set-up. Ellis is informally advising Boakye as he works on his PhD in Electrical Engineering at UC Berkeley. Ellis enjoys working with Boakye, as they have very similar research interests, and hopes to collaborate with him on future projects.

featured alum: eric fosler-lussier



ICSI Alum Eric Fosler-Lussier

Eric Fosler-Lussier is an assistant professor at Ohio State University. He worked at ICSI while he was an Electrical Engineering graduate student at UC Berkeley from 1994-1999, then as a postdoctoral researcher from 1999-2000. He left ICSI after being offered a position at Bell Labs/Lucent Technology, where he worked on development of dialog systems. When Lucent closed their speech group in 2002, he did a short research visit to Columbia University where he worked with our other featured alum Dan Ellis. He then secured an assistant professorship at Ohio State, where he has been since.

His research focus today is on the use of neural net detectors for phones (the basic units of pronunciation used in automatic speech recognition) and features of phones that can be used to improve ASR accuracy. He is looking at the relationship between the phone predicted by a neural net detector versus the dictionary pronunciation, and rethinking how ASR could work in light of this. This current work builds on ASR systems developed at ICSI: Nelson Morgan and Hervé Bourlard's original hybrid system, and the system that followed it, developed by Dan Ellis and Hynek Hermansky, a Gaussian based system using neural nets. Fosler-Lussier is now using feature detectors into conditional random fields (a slightly different statistical paradigm).

While at ICSI, Fosler-Lussier worked for the Speech Group on projects related to signal processing for ASR. In one project, he worked with Professor Nelson Morgan and Dr. Nikki Mirghafori on analyzing speaking rate. At the time, Mirghafori, now a staff scientist with the Speech Group, was writing her master's thesis on the effect of speaking rate on ASR. In an early project, they looked at the wall street journal corpus and the effect of speaking rate on word error rate (WER). Later, they used the Switchboard corpus, which was hand-transcribed at the word, syllable, and phoneme level, using the syllable count from the transcription to determine speaking rate. Morgan led the development of a signal processing method based on this work, and Fosler-Lussier did significant work on that system. They evaluated the output of signal processing versus the manual transcription of switchboard, and used it to predict reduced pronunciations (i.e., slurring syllables together) with regard to speaking rate. They also compared pronunciations from human transcripts to the dictionary pronunciations of the same words as well as to the pronunciation from the signal processing output, again with regard to syllable count and reductions. Eric also briefly worked on the SmartKom project (predecessor to the current SmartWeb project) at ICSI.

As a graduate student, he was a teaching assistant for an artificial intelligence class taught by ICSI emeritus board member Professor Jitendra Malik. For this course, Fosler-Lussier gave a lecture on Hidden Markov Models (HMMs), which was developed into an ICSI technical report, and has since been used as course material since then at several universities.

LATE-BREAKING ALUMNI NEWS

Since conducting our interview with Dan Ellis, we learned that his family has grown to include a new baby girl, Zoe, born on February 6th. Congratulations to Dan and his family!

Jeff Bilmes and Katrin Kirchhoff, both Speech Group alums, also have a new baby. Their son Alexander was born on March 2nd. Jeff and Katrin, who met while working at ICSI, are now on the faculty of University of Washington, where they continue to pursue speech-related research. Congratulations to Katrin and Jeff!

news briefs

Continued from Page 3

ICSI's **BOARD OF TRUSTEES** met on October 13th, and voted on changes to Board membership and corporate officers. Professor **STUART RUSSELL** replaces Jitendra Malik as ex officio board representative from the University of California, Berkeley's Computer Science department, and Professor **JAVIER ARACIL** of Spain's Ministry of Education and Science (MEC) and the Universidad Autonoma de Madrid fills the board seat of recently deceased Dr. Pedro Lizcano. Dr. **JORDAN COHEN** left the Board to focus on his new position as a Senior Scientist at SRI International, a close ICSI collaborator. Replacing him as Secretary-Treasurer is ICSI's Professor **SCOTT SHENKER**, formerly the Vice President of the corporation. Dr. **MARCIA BUSH**, Associate Director of ICSI, is the new Vice President, and Professor **NELSON MORGAN**, Director of ICSI, remains the President of the corporation.

ICSI hosted its annual **RESEARCH REVIEW** on October 13th, along with a book signing party honoring Professor **JERRY FELDMAN**'s new book, **FROM MOLECULE TO METAPHOR**. Professor Feldman was also honored by ICSI's Board of Trustees with the 2006 Distinguished Service Award during the Board of Trustees Meeting that morning. He was recognized for the guidance and expertise he provided during 10 years as ICSI's first Director, and his continued service to the Institute as a scientist in the AI Group. During the Research Review, Professor **SCOTT SHENKER** and Dr. **DILEK HAKKANI-TUR** of ICSI's research staff presented their latest work. Shenker, recipient of the IEEE Internet Award, is the head of the Networking Group. Hakkani-Tur recently joined ICSI from AT&T Research and works on spoken language processing.



Professor Jerry Feldman talks with Professor Wolfgang Wahlster at the research review and book signing following ICSI's Board of Trustees meeting

Professor **DOMENICO FERRARI** (ICSI alum featured in the March 2006 ICSI Gazette) received the 2006 SIGCOMM award. The SIGCOMM award recognizes lifetime contribution to the field of communication networks. Ferrari was selected for "fundamental contributions to Internet QoS architecture and wide-area network research testbeds, his leadership in managing the BSD Unix project, his leadership in founding and establishing ICSI as a premier international networking research institution, and his dedication to educate the next generation of leaders in networking research". A list of past recipients is available on the SIGCOMM web site:
<http://www.sigcomm.org/awards.html>

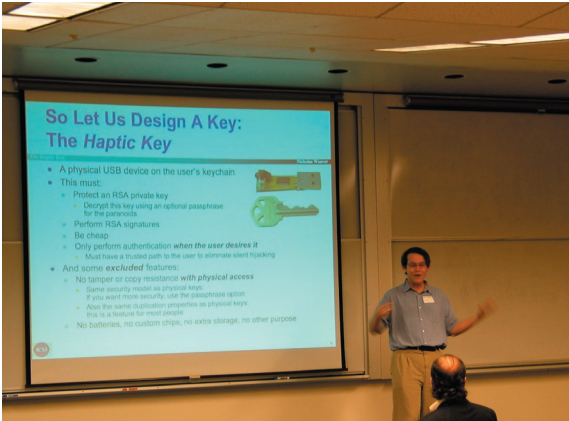
Research by ICSI alum **LISA AZIZ-ZADEH** (now at University of Southern California) on mirror neurons was featured in a September 21, 2006 article on seedmagazine.com. Aziz-Zadeh uses fMRI technology to observe brain activity in subjects when they observe an action and when they read words describing that action. Her

results show that the activity in the brain's premotor cortex is the same for both. Aziz-Zadeh was a postdoctoral researcher with ICSI's NTL project in 2005. Her fMRI work continues to provide physical evidence in support of NTL's hypothesis that mirror neurons play an important role in language.

Professor **VERN PAXSON**, Internet security expert from ICSI's Networking Group, has been awarded the "SIGCOMM Test of Time" award for his paper titled "End-to-End Routing Behavior in the Internet," which was published at SIGCOMM's 1996 conference. The award recognizes a paper published 10 to 12 years in the past, the contents of which continue to be useful today.

Congratulations to Dr. **EDDIE KOHLER**, an ICSI alum now teaching at UCLA, on being selected by MIT Technology Review as a TR35 2006 Young Innovator. Kohler was selected for developing Asbestos, an operating system that protects private data, even when software programs may have failed. Read more about Kohler and his work online at MIT Technology Review. <http://www.technologyreview.com/TR35/index.aspx?year=2006&tchannel=All>

events at ICSI: BEARS Open House



Dr. Nicholas Weaver presents his Haptic Key approach to user authentication at ICSI's Open House

ICSI's annual Open House, in conjunction with the UC Berkeley BEARS conference, was on February 15th from 2:00-4:00. The event featured a talk by Dr. Nicholas Weaver of the Networking Group on his unique approach to User Authentication. He proposes using physical keys, called haptic keys, for each user as an alternative to the current password system, which has many security vulnerabilities.

ICSI researchers presented demonstrations of some of the latest technology developments at ICSI.

Alberto Amengual, a visiting scientist from Spain, presented his computational model of early childhood attachment behavior. The model uses a visual robotic display to show the predicted behavior of a one-year-old child when a stranger enters a room.

Gerald Friedland, a visiting scientist from Germany, showed his diarization system at work parsing the speech of reporters in an Arabic news broadcast. The system is being adapted to work for movie trailers as well, and is being trained to distinguish between speech and other noise typical in movie trailers, such as background music.

Adam Janin, research scientist with the Speech Group, demonstrated the AMIDA Meeting Browser, which takes ASR (automatic speech recognition) output from a recorded meeting, and can automatically browse the text for content.

Luke Gottlieb, also of the Speech Group, presented an interactive demo of the English component to SmartWeb. Users spoke into a microphone, asking questions related to World Cup Soccer, and the SmartWeb system looked up the answers.

FrameNet staff showed some of the latest additions to the FrameNet lexicon, and discussed uses for the annotated frame semantics database. Visiting scientists Kyoko Ohara and Carlos Subirats, from Japan and Spain respectively, talked about efforts to create FrameNet databases in their native languages.

Concurrent with the technology demos, scientists presented posters related to research in progress at ICSI. Nick Weaver presented a poster on network security research, Bonnie Kirkpatrick presented a poster on the latest developments in haplotyping work at ICSI, and a guest poster on reducing network energy consumption was presented by Bruce Nordman of Lawrence Berkeley Laboratory (LBL). Bruce is collaborating with Vern Paxson and Scott Shenker of the Networking Group on a new research proposal.



Luke Gottlieb explains the SmartWeb system to visitors at the open house

bfoit summer institute 2007

As a part of the Information Technology Leadership Program (ITLP) led by the Berkeley Foundation for Opportunities in Information Technology (BFOIT), the annual Summer Institute brings twenty-five underrepresented and disadvantaged high school students with interest and skills in computer science to UC Berkeley for two weeks. BFOIT, which is hosted at ICSI, is currently accepting applications for the 2007 Summer Institute. This year's program takes place from August 6th through the 17th.

There, the students will get hands-on instruction in the fundamentals of computer programming, attend academic and professional workshops, interact with faculty and high tech professionals, and become more familiar with a university environment. In addition to gaining tangible skills, the students' accomplishments and exposure instill a belief in themselves and their ability to pursue a high tech education at a competitive university. Students are also taken on trips to local technology firms. Guest speakers at the Summer Institute last year included Fiona Doyle, the acting Dean of Engineering at UC Berkeley, and ICSI graduate student researcher Kofi Boakye, who is currently working toward a PhD in electrical engineering at UC Berkeley. Kofi's speech recognition demo was a favorite of many BFOIT students.

BFOIT also hosts a summer camp for middle school students, called Science For Youth, or Sci-FY. These programs enable BFOIT to engage with a

cohort of young students, with whom we maintain communication for the Information Technology Leadership Program in the Fall. Students and parents receive targeted advising and coaching on issues like college application, financial aid, and career paths. BFOIT offers scholarships for SAT and ACT preparation courses to its participants.

BFOIT continues to support its graduates during the college application process and throughout their college years with mentoring, scholarship support, and internships.

Program graduates have gone on to degree programs at competitive schools like UC Berkeley and Brown. Students from last year's program are beginning to receive college acceptances from universities nationwide, including UC Davis and Loyola Marymount, where two participants just received full tuition scholarships.

Applications to BFOIT's Summer 2007
Institute will be accepted through
July 20th, 2007.
Eligible students can apply online at
www.bfoit.org

Eligible students interested in the middle school and high school summer programs can download an application at www.bfoit.org. The deadline to apply is July 20, 2007.

BFOIT is also seeking volunteers to assist staff at both the high school and middle school programs. Please see www.bfoit.org for more information.

visiting scholars

Since its inception, ICSI has had a strong international program consisting primarily of ties with specific countries. Current formal agreements exist with Finland, Germany, Spain, and Switzerland.

FROM FINLAND

Jyri Kivinen (Algorithms)
Teemu Koponen (Networking)
Erno Lindfors (Algorithms)
Janne Lindqvist (Networking)
Pasi Sarolahti (Networking)
Antti Vehkaoja (Campus Affiliation)
Mari Zakrzewski (Campus Affiliation)

FROM GERMANY

Gerald Friedland (Speech)
Tobias Kiesling (Networking)
Christian Kreibich (Networking)
Birte Lönneker-Rodman (AI-FrameNet)
Christian Müller (Speech)

FROM SPAIN

Alberto Amengual (AI)
Juan Botia (Algorithms)
Lucia Conde (Algorithms)
Pedro Ruiz (Networking)
Alberto Suarez (Algorithms)
Carlos Subirats (AI-FrameNet)

FROM SWITZERLAND (IM2)

Sebastien Cuendet (Speech)
Matthew Magimai-Doss (Speech)

AMIDA (EUROPEAN UNION)

Joan Biel (Speech)
Marijn Huijbregts (Speech)
Jachym Kolar (Speech)
Beatriz Trueba (Speech)

In addition, we often have visitors associated with specific research and projects.

NETWORKING

Ravi Dhillon
Larry Kao
Chih-Hung Lin
Junda Liu
Vijay Ramachandran
Yun Shi
Shutang Yang

SPEECH

Umit Guz

FRAMENet

Kyoko Ohara
Masayo Kanno
Hiroaki Sato

AI

Artjom Klein

EXTREME ARCHITECTURE

Krste Asanovic

ALGORITHMS

Gad Kimmel
Sriram Sankararaman



Gerald Friedland



Lucia Conde



Jyri Kivinen



Gad Kimmel

publications listing

- M. ALLMAN, E. BLANTON, V. PAXSON, AND S. SHENKER. Fighting Coordinated Attackers with Cross-Organizational Information on Sharing. *Proceedings of HotNets 2006*, November 2006.
- G. ANDEANI, D. DI FABBRIZIO, M. GILBERT, D. GILICK, D. HAKKANI-TUR, AND O. LEMON. Let's DISCOH: Collecting an Annotated Open Corpus with Dialogue Acts and Reward Signals for Natural Language Helpdesks. *IEEE 2006 Workshop on Spoken Language Technology (SLT)*, Aruba Marriot, December 2006.
- U.M. BORGHOFF, P. ROEDIG, J. SCHEFFCZYK, AND L. SCHMITZ. Long-Term Preservation of Digital Documents: Principles and Practices. *Springer-Verlag Berlin Heidelberg*, 2006.
- K. CHAUDHURI, E. HALPERIN, S. RAO, AND S. ZHOU. A Rigorous Analysis of Population Stratification with Limited Data. *Proceedings of ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2007.
- D. CHU, L. POPA, A. TAVAKOLI, J. HELLERSTEIN, P. LEVIS, S. SHENKER, AND I. STOICA. The Design and Implementation of a Declarative Sensor Network System. *Technical Report EECS-2006-132*, EECS Department, University of California, Berkeley, 2006.
- S. CUENDET, D. HAKKANI-TUR, AND G. TUR. Model Adaptation for Sentence Segmentation from Speech. *IEEE 2006 Workshop on Spoken Language Technology (SLT)*, Aruba Marriot, December 2006.
- C. DASKALAKIS, A.G. DIMAKIS, R.M. KARP, AND M.J. WAINWRIGHT. Probabilistic Analysis of Linear Programming Decoding. *Proceedings of ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2007.
- G.V. DRIVONIKOU, P. KAY, T. REGIER, R.B. IVRY, A.L. GILBERT, A. FRANKLIN, AND I.R.L. DAVIES. Further evidence that Whorlian effects are stronger in the right visual field than the left. *Proceedings of the National Academy of Sciences*. 104, 1097-1102. (2007).
- C. FILLMORE, S. NARAYANAN, AND C. BAKER. What Linguistics can Contribute to Event Extraction. *AAAI Workshop on Event Extraction, AAAI 2006*, Boston, 2006.
- S. FLOYD. RFC 4774: Specifying Alternate Semantics for the Explicit Congestion Notification (ECN) Field. *Request For Comments 4774*, Best Current Practice, November 2006.
- S. FLOYD, M. ALLMAN, A. JAIN, AND P. SAROLAHTI. RFC 4782: Quick-Start for TCP and IP. *Request For Comments 4782*, Experimental, January 2007.
- B. GODFREY AND R.M. KARP. On the Price of Heterogeneity. *Parallel Systems Symposium on Parallel Algorithms and Architectures*, 2006.
- E. HALPERIN, G. KORTSARZ, R. KRAUTHGAMER, A. SRINIVASAN, AND N. WANG. Integrality ratio for Group Steiner Trees and Directed Steiner Trees, *SIAM Journal on Computing*, Vol. 36, 1494-1511, 2007.
- T. HENDERSON, S. ROY, S. FLOYD, AND G. RILEY. NS-3 Project Goals. *Proceedings of WNS2 (Workshop on ns-2: the IP Network Simulator)*, October 2006.
- D. HILLARD, Z. HUANG, H. JI, R. GRISHMAN, D. HAKKANI-TUR, M. HARPER, M. OSTENDORF, AND W. WANG. Impact of Automatic Comma Prediction on POS/Name Tagging of Speech. *IEEE 2006 Workshop on Spoken Language Technology (SLT)*, Aruba Marriot, December 2006.
- A. JANIN, A. STOLCKE, X. ANGUERA, K. BOAKYE, O. CETIN, J. FRANKEL, AND J. ZHENG. The ICSI-SRI Spring 2006 Meeting Recognition System. In S. Renals and S. Bengio, editors, *Machine Learning for Multimodal Interaction: Third International Workshop (MLMI 2006)*, Lecture Notes in Computer Science. Springer, 2007.
- J. KANNAN, J. JUNG, V. PAXSON, AND C. KOKSAL. Semi-Automated Discovery of Application Session Structure. *Proceedings of ACM Internet Measurement Conference*, October 2006.
- R.M. KARP. Fair Bandwidth Allocation Without Per-Flow State. *Essays in Memory of Shimon Even*, Springer-Verlag, 88-110, 2006.
- R.M. KARP. George Dantzig's Impact on the Theory of Computation. To appear in *Discrete Optimization*, 2007.
- R.M. KARP AND R. KLEINBERG. Noisy Binary Search. *Proceedings of ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 881-890, 2007.
- P. KAY AND T. REGIER. Color naming universals: the case of Berinmo. *Cognition*. 2007 Feb, 102(2): 289-98. 2007.
- K. LAKSHMINARAYANAN, A. RAO, I. STOICA, AND S. SHENKER. End-host Controlled Multicast Routing. *Elsevier Computer Networks, Special Issue on Overlay Distribution Structures and their Applications*, 2006.
- S. KIM, R. FONSECA, P.K. DUTTA, A. TAVAKOLI, D.E. CULLER, P. LEVIS, S. SHENKER, AND I. STOICA. Flush: A Reliable Bulk Transport Protocol for Multihop Wireless Network. *Technical Report EECS-2006-169*, EECS Department, University of California, Berkeley, December 2006.
- E. KOHLER, J. LI, V. PAXSON, AND S. SHENKER. Observed Structure of Addresses in IP Traffic. *IEEE/ACM Transactions on Networking*, Vol. 14, No. 6, December 2006.
- D. LIU, M. ALLMAN, S. JIN, AND L. WANG. Congestion Control Without a Startup Phase. *Protocols for Fast, Long Distance Networks (PFLDnet) Workshop*, February 2007.
- Y. LIU, N.V. CHAWLA, M.P. HARPER, E. SHRIBERG, AND A. STOLCKE. A study in machine learning from imbalanced data for sentence boundary detection in speech. *Computer Speech and Language*, vol. 20, pp. 468-494, October 2006.
- R. PANG, V. PAXSON, R. SOMMER, AND L. PETERSON. binpac: A yacc for Writing Application Protocol Parsers. *Proceedings of ACM Internet Measurement Conference*, October 2006.
- A. PARKER, S. REDDY, T. SCHMID, K. CHANG, G. SAURABH, M. SRIVASTAVA, M. HANSEN, J. BURKE, D. ESTRIN, M. ALLMAN, AND V. PAXSON. Network System Challenges in Selective Sharing and Verification for Personal, Social, and Urban-Scale Sensing Applications. *Proceedings of HotNets 2006*, November 2006.
- T. REGIER, P. KAY, AND N. KHETARPAL. Color naming reflects optimal partitions of color space. *Proceedings of the National Academy of Sciences*. 104, 1436-1441. (2007).
- G. TUR, U. GUZ, AND D. HAKKANI-TUR. Model Adaptation for Dialog Act Tagging. *IEEE 2006 Workshop on Spoken Language Technology (SLT)*, Aruba Marriot, December 2006.
- N. WEAVER, V. PAXSON, AND J.M. GONZALEZ. The Shunt: An FPGA-Based Accelerator for Network Intrusion Prevention. *ACM/SIGDA International Symposium on Field Programmable Gate Arrays (FPGA)*, February 2007.
- M.A. WEBSTER AND P. KAY. Individual and Population Differences in Focal Colors. *The Anthropology of Color*, R.E. MacLauray, G.V. Paramei, D. Dedrick, and J. Benjamins, eds., 2006.

featured research, con't.

significantly improved word error rate (WER) for both Mandarin and Arabic in initial tests performed at University of Washington and SRI. Encouraged by the early positive results, scientists began to work on incorporating new acoustic features, which are already providing improvements in WER. The new approach may be adapted for use with SmartWeb, another collaborative speech project involving ICSI scientists. (More information on SmartWeb is available at <http://www.smartweb-project.de>.)

DIARIZATION

When analyzing the output of ASR systems, one challenge is to determine who is speaking and when a new speaker enters the conversation. ICSI researchers have developed highly successful methods for automatically parsing the speech signal based on who speaks when - this is called diarization. The ICSI diarization system was originally developed as part of the ongoing effort in processing speech from meetings, which are all in English. For the GALE project, researchers processed 696 hours of new data in English, Mandarin, and Arabic.

SENTENCE SEGMENTATION

Sentence segmentation refers to the automatic segmentation of a stream of words into sentences. This makes the output of an ASR system much more comprehensible to a human, but also is crucial for other automatic processing tasks of the speech signal, such as machine translation and information extraction.

INFORMATION DISTILLATION

Information distillation is the extraction of the most useful pieces of information related to a given query from massive multilingual audio and text documents. As an example, if someone needed information about

an individual, and had some audio and text files to search through, a system that performs information distillation could be very useful. Instead of manually scanning through text documents and listening to the audio files, which is extremely time consuming for people, the distillation system could find relevant information from the files in seconds. Distillation work at ICSI, led by Dilek Hakkani-Tur, began in May and produced positive results as early as June in the first round of NIST evaluations.

OTHER SPEECH PROCESSING RESEARCH AT ICSI

In addition to the GALE effort, the Speech Group continues to work on ASR for meetings, dialog systems, and speaker recognition, as well as the Tamil speech recognizer, a project bringing access to information technology (via a voice controlled user interface) to remote villages in India. Most of the visiting scientists to the Speech Group participate through two international collaborative projects focused on meetings, Interactive Multimodal Information Management (IM2) and Augmented Multi-party Interaction Distance Access (AMIDA, the successor to AMI).

To learn more about current research by the Speech Group, browse the group's publications at <http://www.icsi.berkeley.edu/cgi-bin/pubs/index.pl> (select "Speech" from the drop down menu) or read the Speech section of the 2006 Annual Report, which can be downloaded in .pdf format (select "ICSI" from the drop down menu on the publications page). The Annual Report contains detailed technical information on all areas of research at ICSI from the past year.



INTERNATIONAL
COMPUTER SCIENCE
INSTITUTE

THE INTERNATIONAL COMPUTER SCIENCE INSTITUTE (ICSI) is the only independent, non-profit US lab solely conducting open, non-proprietary, pre-competitive research in computer science. Affiliated with the University of California campus in Berkeley, ICSI's mission is to further research in computer science through international collaboration, and further international collaboration through research in computer science. ICSI provides a haven for computer scientists to conduct long-term research without commercial limitations and with few faculty pressures. ICSI has significant efforts in four major research areas: Internet research, including Internet architecture, related theoretical questions, and network services and applications; theoretical computer science, including applications to bioinformatics;

artificial intelligence, particularly for applications to natural language understanding; and natural speech processing.

BOARD OF TRUSTEES

Javier Aracil
Hervé Boulard
Beth Burnside
Adele Goldberg
Greg Heinzinger
Clifford Higgerson
Richard Karp
Nelson Morgan
David Nagel
Stuart Russell
Jouko Salo
Shankar Sastry
Scott Shenker
Wolfgang Wahlster

The ICSI Gazette is published twice a year by the International Computer Science Institute (ICSI). All material in *The ICSI Gazette* is copyrighted by ICSI and should not be copied without explicit prior written permission.

EDITOR-IN-CHIEF

Nelson Morgan

SENIOR EDITOR

Maria Quintana

DESIGN

envision ID

LAYOUT

Leah Hitchcock

COPYWRITER

Leah Hitchcock

PHOTOGRAPHY

Diane Starr
Leah Hitchcock

INTERNATIONAL
COMPUTER
SCIENCE
INSTITUTE

1947 Center Street Suite 600
Berkeley, CA 94704 USA
v 510 666 2900
www.icsi.berkeley.edu