

## featured icsi research: multimedia

A team of researchers spanning the Speech, Vision, and Networking groups is collaborating on ways to extract meaning from the vast amounts of multimedia data freely available online - a dataset that includes hours of new videos and thousands of new photos uploaded to the Internet every minute. This unconstrained data - videos and images unregulated for quality, size, or content - presents challenges for techniques known to recognize sounds and images successfully in laboratory conditions. But by analyzing multiple modalities - by combining, for example, audio processing techniques with image analysis - the researchers hope to design methods for creating meaning out of the enormous amounts of data available.

"The only way to solve these problems is to be as open as the dataset," said Gerald Friedland, a senior researcher in the Speech Group who leads ICSI's multimedia efforts. "It's about taking everything you can into account."

The team works closely with UC Berkeley's Parallel Computing Laboratory, which provides new algorithmic ideas that help the researchers deal with the huge quantity of data. ParLab is funded by Microsoft and Intel.

ICSI researchers have been working on acoustic recognition since the early 1990s, and on computer vision since 2008. But recent efforts are the first that aim to combine acoustic and visual recognition in order to extract meaning from data.

### MULTIMODAL LOCATION ESTIMATION

Using both modalities, the team is designing ways to identify automatically where consumer-produced videos were taken. The team has applied ICSI's speaker diarization system to audio tracks taken from a dataset of unconstrained videos, building audio profiles of the 18 cities where the videos were taken. The speaker diarization system is traditionally used on audio tracks

containing human speech; it identifies who spoke when. But the team found that the system can identify where videos were shot by detecting subtle cues even in audio tracks that do not contain human speech.

"It's hard for humans to do this task," said Howard Lei of the Speech Group, "but our algorithms were able to pick things up" that a human might not. Future work, to be supported by the National Science Foundation, will look at how accurately humans are able to place the videos.



When the researchers discounted videos submitted by the same user, the system correctly identified almost 68 percent of the videos. When videos from the same users are included, the system identified almost 80 percent of the videos correctly.

The work is supported by the National Geospatial-Intelligence Agency University Research Initiatives Program.

It's unclear exactly what sounds the system uses to identify cities. While the researchers noticed some audio trends specific to certain cities - for example, the sound of birds is prominent in videos from San Francisco, while the sound of trains is prominent in videos from Tokyo - they "still don't know which features of the audio this system picks up on," said Jaeyoung Choi of the Speech Group.

Choi has begun to add visual recognition techniques to the system to improve accuracy. "Visual [analysis]

### in this issue:

page 2  
As I See It  
by Nelson Morgan, Director

page 3  
News Briefs

page 6  
Netalyzr

page 7  
Visiting Scholars

page 9  
Publications



## as i see it by Nelson Morgan, Director

*Plus ça change, plus c'est la même chose.*

- *Jean-Baptiste Alphonse Karr*

ICSI is about to change – as it has many times in its quarter-century of existence. But indeed, it will stay the same.

And this time, the change (and also the continuation) is quite personal. I'm leaving the position of ICSI director at the end of the year, after almost 13 years in this position. However, I'm not leaving ICSI. My current plan is to continue at the Institute, mostly as a research scientist, although I will be giving some assistance to the new director. ICSI has been a wonderful place for me, and I have no wish to leave it. But 13 years at the helm is probably enough. And as previous director Jerry Feldman told me when I took the job, "There are two great days in the life of a boat owner: the day he gets the boat, and the day that he lets it go."

I came here as a research scientist in 1988, and within a few years ICSI was not just another job; it was my home. So when the director opportunity "knocked" in 1999, I enthusiastically accepted. A couple of months later I wasn't so sure that this had been a very wise move; the Institute was transitioning between two different funding models, and it was more than a little scary while we figured out how to stay afloat. But with the help of our outstanding staff, we made it through a tough time, and before long we had grown into a stronger research organization. We've had to change our funding model as conditions changed. We moved from a reliance on international sources to a budget largely based on industrial funding, during the dot-com boom; and then, with the dot-com bust, our primary support came from U.S. Federal sources, even as significant international and industrial funding continued. The cast of characters at ICSI has also changed quite a bit over the years; of the group that was here for ICSI's formal inauguration in 1988, only Jerry Feldman, Dick Karp, and I are still on staff. But despite changes in research directions, funding model, and personnel, in many ways ICSI seems to be pretty much the same place. We still have our "teas" every Tuesday and Thursday afternoons, a practice that Jerry started early on. We still hear multiple languages in the halls. We

still have an extremely responsive administrative staff, one that handles the heterogeneous needs of our researchers with skill and grace. The accomplishments of ICSI principal investigators have certainly not decreased; to give just one example, four out of the last ten SIGCOMM Award winners have come from ICSI (including, this year, Vern Paxson). And of course, we still have our close relationship with the UC Berkeley campus, especially with EECS, while maintaining the independence that keeps us flexible in so many ways.

And so, ICSI largely stays the same, despite continual change. And once again, ICSI is about to change, though not in any way that should diminish it.

Starting in January, our new director will be Dr. Roberto Pieraccini, an experienced and respected scientist who has excelled in the area of spoken dialog systems and other research topics at well known laboratories like AT&T and IBM, and who also has leadership experience in another small independent organization (in his case, a successful commercial start-up). We've spoken extensively about ICSI's future and I'm confident that he will be great in the director role. I'm also going to assist Roberto as his deputy director, at least for a while.

It's been a great privilege to be ICSI's director. The easiest part of the job has been bragging about our amazing staff – I don't have to work very hard to find reasons to praise the efforts of people who are at the top of their field. The administrative staff is also exceptional; it is almost routine to have glowing comments from visitors in their departure documents.

Given the upcoming change, this is my last opinion column for the ICSI Gazette. I've had a lot of fun spouting off twice a year, and pointing to some of the wonderful things about ICSI. I've particularly enjoyed picking out the quotes that I start these columns off with, so I will indulge in a double dip here by closing with this one:

*Happy trails to you, until we meet again ...*

- *Dale Evans*

# icsi announces new director

ICSI is pleased to announce the appointment of Roberto Pieraccini to the position of director. In January, Pieraccini will succeed Nelson Morgan, who will continue to lead the Speech Group and assist Pieraccini as his deputy director. Pieraccini comes to ICSI from SpeechCycle, where he has served as chief technology officer since 2005. Before that, he was the manager of the Conversational Interaction Technology department at the IBM T. J. Watson Research in Yorktown Heights. He has authored more than 130 publications in the areas of speech recognition, spoken language understanding and dialog, multimodal interaction, standards for human-machine communication, and machine learning. He is a fellow of the IEEE and ISCA, a member of the AVIOS board, and a member of the editorial boards of several scientific and technology magazines. He graduated in electrical engineering from the Università degli Studi di Pisa in Pisa, Italy, in 1980, and has worked at CSELT, AT&T Bell Laboratories, AT&T Shannon Laboratories, and SpeechWorks International.

Morgan has directed ICSI since 1999 and has led the Speech Group since ICSI's inauguration.



## news briefs

The 2011 ACM SIGCOMM Award was given to UC Berkeley Professor **VERN PAXSON**, a senior researcher in the Networking Group. Paxson was recognized for his seminal contributions to Internet measurement and security, as well as for distinguished leadership and service to the Internet community. Past winners of the award include UC Berkeley Professor **SCOTT SHENKER**, who leads the Networking Group, **SALLY FLOYD**, a senior researcher in the group, and **DOMENICO FERRARI**, a former deputy director of ICSI and Networking Group leader. Paxson was also invited to give the keynote talk at ACM SIGCOMM 2011, the annual conference of the ACM Special Interest Group on Data Communications. SIGCOMM is considered the leading conference in networking.

The ACM SIGCOMM Test of Time Award was presented to Networking Group researchers for their 2001 paper on scalable content-addressable networks. The award recognizes papers that continue to be useful contributions to the field a decade after publication. The award was given for "A Scalable Content-Addressable Network," written by **SYLVIA RATNASAMY**, now a professor at UC Berkeley and an ICSI external fellow; ICSI alumni **PAUL FRANCIS**, director of the Max Planck Institute for Software Systems, and **MARK HANDLEY**, professor at University College London; and UC Berkeley professors **RICHARD KARP** and **SCOTT SHENKER**, who lead the Networking and Algorithms groups respectively.

Two PhD students in the Speech Group have received a Qualcomm Innovation Fellowship to pursue their work on a system that automatically summarizes large amounts of text. **MOHIT BANSAL** and **TAYLOR BERG-KIRKPATRICK** will collaborate with Qualcomm research and development throughout the year and will receive \$50,000 each in the coming year to work on their project, "Automatic Summarization for Mobile Search." They are designing a system that, for example, can produce a brief synopsis of recent news resulting from an Internet search. The Qualcomm Innovation Fellowship is awarded annually to help PhD students pursue innovative ideas in computer science or electrical engineering. This year, more than 140 teams applied for eight fellowship slots.

*continued on page 8*

# featured research: continued

*continued from page 1*

is more straightforward” than audio analysis, he said. The researchers will analyze visual elements such as which way lines are oriented, which might suggest sea lines or tall buildings. They will also analyze what textures appear in videos, which helps to identify pavement, grass, and other features, and how color is distributed.

Choi will use millions of photos pulled from the Internet in order to develop visual profiles for different cities around the world. The researchers will compare a frame from the video that they are trying to place against the cities’ profiles to find which profile matches the frame most closely.

The researchers decided to pursue this technique – called nearest-neighbor matching – rather than trying to detect landmarks unique to particular cities. “In terms of research value, landmark recognition has already been done,” said Choi, while identifying scenes based on visual elements “is more of an emerging field.”

## **VIDEO CONCEPT DETECTION**

Researchers are also working toward a system that can detect concepts in videos – that can, for example, search large collections of videos for those that match statements like “feeding an animal.” In the ALADDIN program, funded by IARPA, teams from institutions around the world are building a concept detection system. IARPA has provided the teams with tens of thousands of consumer-produced videos, some of which are labeled as belonging to one of 15 categories. Given the labeled examples, the challenge is to find videos that belong in any of the 15 categories from a set of about 50,000 unlabelled videos. ICSI researchers, working closely with SRI, Carnegie Mellon University, and other research institutions, are using acoustic analysis in two main efforts toward the goal of detecting concepts in these videos.

The first system involves what Robert Mertens calls “holistic analysis.” Mertens is visiting ICSI from Germany to work full time on the project, which is led by Friedland. The researchers built a model for each category based on all acoustic features from all videos labeled as belonging to that category. “We try to find and learn configurations of audio features that indicate that a video belongs to a category,” Mertens said. The team also tracked those features that occurred frequently in all videos, which aren’t helpful in determining which category a video belongs to. While the process works, it provides little insight into how it works. Mertens said the method is like

“looking into a black box”: researchers aren’t able to say what aspects of a video their system identifies as belonging to a certain category.

Since one goal of the project is to be able to explain how a system gets its results, Mertens and the team also used ICSI’s speaker diarization system to group the videos’ audio tracks into segments of similar sounds, developing a profile for each category based on the sounds that frequently occur in the videos belonging to that category. The speaker diarization system is traditionally used to identify who spoke when in audio tracks that contain speech. The system usually only takes into account segments of sound that last for two seconds or longer, and segments of sound that contain speech. The team modified the system so that not only did it analyze sounds that are not speech, but it also analyzed much shorter segments of sound, such as drum beats. The system then analyzed the videos labeled as belonging to a category and identified the sounds that best represented the category based on how frequently the sounds occurred in each video. The videos without labels are then searched to see which match the profiles. The sounds are essentially treated as words: if, for example, the word “thoracic” appears frequently in a book, it’s plausible the book is medical in nature. Similarly, if a particular sound occurs frequently in a video’s audio track, the video may belong to a category that contains other videos with the same sound. The team found 300 sounds that could be used to identify which category a video belongs to.

Eventually, says Mertens, the method might allow researchers to ignore certain sounds that only incidentally occur in a category, or to ignore certain combinations of sounds. For example, if a user is searching for a video depicting a wedding, it might be helpful to include sounds of applause, but not of guitar music typically associated with hard rock concerts. Mertens’s system may allow users to exclude these combinations. The system is also “more explicable,” said Mertens, “because you can tell the user the system’s decision is based on the frequency of occurrence of these sounds.”

## **PRIVACY WORK**

The multimedia work has serious implications for people who upload videos and photos to Web sites like YouTube and Flickr: if it is possible to identify where videos were taken and what they depict using visual and acoustic recognition methods, it might also be possible to find out a great deal of information about the people who have created them.

Last year, Friedland and Robin Sommer, a senior researcher in the Networking Group, found that they could identify, for example, the home addresses of people who were on vacation by extracting the longitude and latitude embedded in photos posted to Web sites like Craigslist and Flickr. Many high-end smart phones and digital cameras embed geo-tags, precise coordinates showing where a photo or video was taken. By combining geo-tags from different videos – say, one labeled “home” and one labeled “vacation” – the researchers were able to find the home addresses of people currently on extended trips.

With billions of videos publicly accessible on the Web, millions of users could potentially be identified across various online profiles.

Researchers are now turning their attention to what they can learn from analysis of video and audio tracks. In work presented at ICASSP this year and supported by the National Science Foundation, researchers used speaker recognition methods on the audio tracks of Flickr videos to determine whether they

were uploaded by the same user. Lei of the Speech Group said the technique “allows us to tie together the identifications of different profiles. We’re trying to raise this concern.”

Lei, Choi, Adam Janin, and Friedland, all of the Speech Group, trained ICSI’s speaker recognition machine on videos posted to Flickr. In laboratory conditions, with high-quality audio tracks of controlled lengths and content, the machine is good at identifying who is speaking. Consumer-created videos like those uploaded to Flickr, on the other hand, are limited to 90 seconds; of the videos used by ICSI researchers in the ICASSP work, one-third were 20 seconds or shorter. The videos may also include any sounds – passing cars or a neighbor’s music; until recently, research in the field of acoustic processing has not focused on such sounds. ICSI’s multimodal researchers, however, “wanted to see how established approaches deal with the random data we’re getting,” said Lei, and this required the analysis of audio content that lies outside the realms of speech and music, on which much research has been done.

Despite the widely varying quality and short length of the videos, however, researchers were able to identify about 66 percent of users.

The work showed how simple it was to link different online profiles, using just the audio track. “There are so many cues that [Web sites] can’t really protect privacy,” Friedland said. For example, while some Web sites claim to protect users’ anonymity – dating services, for example – videos posted there may be matched to videos posted to more public profiles, such as those on YouTube.

And with billions of videos publicly accessible on the Web, millions of users could potentially be identified across various online profiles.

In other work, Choi and Friedland modified the earlier geo-tagging study, which relied on geo-coordinates extracted from YouTube videos to find home addresses of people on vacation who are potential victims of burglary while out of town. They found they could achieve similar results even without geo-tags. They extracted text from the tags added to videos (such as “San Francisco”) and ran it through several filters to account for ambiguities. These filters, all derived from programs freely available on the Internet, helped prevent inaccuracies arising from tags that could refer to multiple places (for example, Paris, France and Paris, Texas), that contained misspellings or incorrect spacing (sanfrancisco for San Francisco), and that included words that could refer either to the name of a location or to something else entirely (the word “video” in a tag, for example, is probably not referring to Video, Brazil). While the accuracy of this method is low, it does not depend on geo-tags, which are embedded in only a fraction of the videos uploaded to the Internet. This means the method can be applied to a far larger set of videos. The researchers were thus able to positively identify the same number of potential victims as the earlier work did.

## OTHER APPLICATIONS

The multimodal work with unconstrained data has applications beyond analysis of consumer-created content found online. Researchers are also working with a UC Berkeley team to construct a robot that can both see and hear in order to improve its capacity to deal with real-world situations. Friedland is also contributing to the Robust Automatic Transcription of Speech project, funded by DARPA, which seeks to improve the accuracy of speech processing tasks, such as speaker identification, for sources of poor quality. Friedland is developing a system that can identify the parts of an audio track that contain speech and the parts that do not. He is currently working with consumer-produced videos to train the system.

# netalyzr: from diagnostic to research tool

Some Internet service providers (ISPs) redirect their customers' Internet searches through third-party companies, according to Networking Group researchers. The behavior is one of several that researchers have tracked using data gathered by Netalyzr, an online tool that tests how open and transparent a user's connection to the Internet is. Netalyzr has been used more than 350,000 times since it went live in 2009, and this summer received recognition from both the Internet Society and the Federal Communications Commission.

Tests run by Netalyzr have provided data for researchers about broad trends in how people connect to the Internet.

While the test was designed to provide information that individual non-expert users could understand easily, tests run by Netalyzr on users' Internet connections have also provided data for researchers about broad trends in how people connect to the Internet.

"Netalyzr's continuing popularity allows us to track and quickly understand emerging connectivity problems encountered by Internet users," said Networking Group researcher Christian Kreibich.

Using data aggregated from Netalyzr sessions, Kreibich, Nicholas Weaver, and Vern Paxson found that a dozen ISPs with millions of customers redirected some Web searches for terms related to popular brands like Apple and Safeway. The ISPs redirected all user traffic for sites like Bing or Yahoo! to third-party proxy servers. Normally these proxy servers would simply relay user requests. The proxies also monitored user requests in order to detect when the user searched for a specific keyword using the browser's search or address bar. If the search matched a keyword related to certain brands, the proxies would replace the normal search results with a redirection through affiliate marketing companies and ultimately to the related brand's Web sites. ISPs apparently earned commissions off the redirected searches, which bypassed the search engines for which they were intended. Researchers identified more than 160 search terms redirected this way, including "Dell," "Bloomingdales," "WSJ," and "CA," which was redirected to

the Computer Associates' Web store. It is unclear whether the affiliate programs or companies holding the brands were aware of this behavior.

The behavior is an extension of another trend the researchers also tracked: users who misspell domain names in their browser bars are sometimes taken to an error page filled with ads, rather than a page supplied by the browser with a message that the server named in the request does not exist. It is believed that ISPs get a cut of the money paid to the marketing companies that generate the ads.

The finding that ISPs were redirecting not just misspelled domain names but also Web searches was reported in the New Scientist on August 4. The next day, the system was featured on popular technology news blog Slashdot.org. Subsequent demands for the Netalyzr test jumped from about 700 requests per day to about 700 per hour.

Also on August 5, the Federal Communications Commission announced that the Netalyzr team had won the FCC Open Internet Research Challenge. FCC Chairman Julius Genachowski said that Netalyzr, along with the other winners of the challenge, is "giving consumers and researchers the information they need to understand and monitor the free and open Internet."

"Netalyzr's continuing popularity allows us to track and quickly understand emerging connectivity problems encountered by Internet users."

- Christian Kreibich  
Networking Group Researcher

Netalyzr is also one of four end-user tests recommended by the Internet Society that determine what problems users may encounter with IPv6, the new Internet protocol address standard that will eventually replace the protocol used by most Internet services today.

Netalyzr is primarily supported by the National Science Foundation, with additional support from Amazon.com, Google, and Comcast.

# visiting scholars

Since its inception, ICSI has had a strong international program consisting primarily of ties with specific countries. Current formal agreements exist with Brazil, Finland, Germany, and Switzerland. In addition, we often have visitors associated with specific research and projects.

## AI

Sergio Guadarrama  
Frank Hopfgartner (Germany)  
Emanuel Kitzelmann (Germany)  
Ru Li  
Jesús Oliva  
Hiroaki Sato  
Malte Schilling (Germany)  
Carlos Subirats  
Tiago Torrent (Brazil)

## ALGORITHMS

Karthekeyan Chandrasekaran  
Shuai Cheng Li  
Matthias Mnich (Germany)  
Roded Sharan

## SPEECH

Arlo Faria (Brazil)  
Joaquin Gonzalez  
Bernd Meyer (Germany)  
Gerald Penn  
Stefan Steidl (Germany)

## ARCHITECTURE

Paula Herber (Germany)  
Miquel Moretó Planas

## NETWORKING

Greg Bodwin  
Jan Böttger  
Gilbert Clark  
Lorenzo DeCarli  
Mohan Dhawan  
Haixin Duan  
Oana Goga  
Dorgival Guedes (Brazil)  
Junaid Khalid  
Gregor Maier (Germany)  
Syed Akbar Mehdi  
Abdelaziz Mohaisen  
Christian Rossow  
Ankit Singla  
Renata Teixeira  
Amin Tootoonchian  
Jack Wu  
Baohua Yang

## VISION

Peer Stelldinger (Germany)

## CAMPUS AFFILIATION / OTHER

Christoph Goebel (Germany)  
Eduardo Huerta (Brazil)  
Kai Huotari (Finland)  
Tommi Lampikoski (Finland)  
Nils Peters (Germany)



Eduardo Huerta



Bernd Meyer



Renata Teixeira




Syed Akbar Mehdi

## news briefs, cont.

continued from page 3

A second edition of the acclaimed textbook *Speech and Audio Signal Processing* by **NELSON MORGAN**, Director of ICSI, and Ben Gold was released in August by Wiley Publishing. The textbook, initially published in 1999, was updated and expanded in a number of areas, including psychoacoustic audio coding and music transcription. Columbia University Professor and ICSI contributor **DAN ELLIS** is a new principal coauthor and editor for the edition. Other new contributors include ICSI researchers **STEVEN WEGMANN** and **GERALD FRIEDLAND** and ICSI alum **DAVID VAN LEEUWEN**.



The *Proceedings of the National Academy of Sciences* has published two studies coauthored by AI researcher **PAUL KAY** on interactions between color, language, and brain activity. One article, published in April, shows that learning new color names causes rapid growth in grey matter in the adult brain. Another article, published in August, shows that the brain-localized effect of linguistic categories on early color processing occurs out of awareness. Kay has published a dozen papers in PNAS in the last several years.

Networking Group researcher **BARATH RAGHAVAN** won the best paper award at the Workshop on Green Networking at ACM SIGCOMM 2011, held August 19 in Toronto. His paper, written with UC Berkeley postdoctoral fellow Justin Ma, explores what might happen to the Internet in the event of an oil shortage resulting in a global energy crisis.

Networking Group researchers have found that the majority of the most active malware distributors pay third parties to install their malicious software on at least some of the computers they infect. One of two Outstanding Paper Awards at the USENIX Security Symposium 2011 was awarded to ICSI's **CHRIS GRIER**, **CHRISTIAN KREIBICH**, and Professor **VERN PAXSON** and

Juan Cabellero of IMDEA Software Institute for the work, which was featured in the MIT Technology Review. The researchers infiltrated four “pay-per-install” providers and downloaded over a million instances of malware. They found that twelve of the 20 most frequently seen families of malware distributors

hire pay-per-install providers to infect machines, spending just pennies per machine infected. The findings suggest that even if a botnet – a potentially very large group of malware-infected computers under the unified command of a single person – is completely wiped out, it could be inexpensively rebuilt from scratch.

Several ICSI researchers will participate in the recently announced Intel Science and Technology Center for Secure Computing at UC Berkeley. Professors **SCOTT SHENKER**, **VERN PAXSON**, and **SYLVIA RATNASAMY** will join researchers from across the U.S. in exploring ways to guard personal computers from malware, secure mobile devices that use third-party applications, and protect personal data on the Internet. The center will receive \$2.5 million per year in funding from Intel.

In a study featured in the New York Times, ICSI, UC Berkeley, and UC San Diego researchers found that just three banks authorize 95 percent of credit card sales of goods advertised through spam. The study measured the spam-based business cycle from the sending of unwanted email to the delivery of goods. **CHRISTIAN KREIBICH**, **NICHOLAS WEAVER**, and Professor **VERN PAXSON**, with researchers at UC San Diego and UC Berkeley, visited 6 million Web sites advertised through spam and purchased more than 100 items such as over-the-counter medications and replica goods in order to understand the economy of spam-based sales. The research suggests profits resulting from these sales could be significantly reduced if credit card-issuing banks refused to settle transactions authorized by banks identified as spam business supporters. The findings were presented at the IEEE Symposium on Security and Privacy. The researchers' work on the profitability of spam has also been featured in *U.S. News and World Report*, *Wired Magazine*, and the NPR News Blog.

ICSI recently welcomed two new arrivals: Virginia Anne, daughter of Networking Group postdoctoral fellow **GREGOR MAIER** and his wife Catherine Logan, born March 21; and Jonathan Haruto, son of Systems Administrator **ADAM RICHMAN** and his wife Akiko, born June 5.



Jonathan Haruto Richman



# publications listing

- M. AKBACAK, D. VERGYRI, A. STOLCKE, N. SCHEFFER, AND A. MANDAL. *Effective Arabic Dialect Classification Using Diverse Phonotactic Models*. Proceedings of the 12th Annual Conference of the International Speech Communication Association (Interspeech 2011), Florence, Italy, pp. 737-740, August 2011.
- G. ANANTHARAYANAN, A. GHODSI, S. SHENKER, AND I. STOICA. *Disk-Locality in Datacenter Computing Considered Irrelevant*. Proceedings of the 13th Workshop on Hot Topics in Operating Systems (HotOS XIII), Napa, California, May 2011.
- S. ARIANFAR, T. KOPONEN, B. RAGHAVAN, AND S. SHENKER. *On Preserving Privacy in Information-Centric Networks*. Proceedings of the ACM SIGCOMM Workshop on Information-Centric Networking (ICN-2011), Toronto, Canada, pp. 19-24, August 2011.
- A. BATLINER, S. STEIDL, AND E. NÖTH. *Associating Children's Non-Verbal and Verbal Behaviour: Body Movements, Emotions, and Laughter in a Human-Robot Interaction*. Proceedings of the IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP 2011), Prague, Czech Republic, pp. 22-27, May 2011.
- Y. BEN-DAVID, S. HASAN, J. PAL, M. VALLENTIN, S. PANJWANI, P. GUTHEIM, J. CHEN, AND E. BREWER. *Computing Security in the Developing World: A Case for Multidisciplinary Research*. Proceedings of the Fifth ACM Workshop on Networked Systems for Developing Regions (NSDR), Washington, D.C., pp. 39-44, June 2011.
- K. BOAKYE, O. VINYALS, AND G. FRIEDLAND. *Improved Overlapped Speech Handling for Speaker Diarization*. Proceedings of the 12th Annual Conference of the International Speech Communication Association (Interspeech 2011), Florence, Italy, pp. 941-944, August 2011.
- J. CABELLERO, C. GRIER, C. KREIBICH, AND V. PAXSON. *Measuring Pay-per-Install: The Commoditization of Malware Distribution*. Proceedings of the 20th USENIX Security Symposium (Security '11), San Francisco, California, August 2011.
- K. CHANG, H. LEI, AND J. CANNY. *Improved Classification of Speaking Styles for Mental Health Monitoring using Phoneme Dynamics*. Proceedings of the 12th Annual Conference of the International Speech Communication Association (Interspeech 2011), Florence, Italy, pp. 85-88, August 2011.
- J. COLMENARES, I. SAXTON, E. BATTENBERG, R. AVIZIENIS, N. PETERS, K. ASANOVIC, J. KUBIATOWICZ, AND D. WESSEL. *Real-Time Musical Applications on an Experimental Operating System for Multi-Core Processors*. Proceedings of the International Computer Music Conference (ICMC 2011), Huddersfield, England, July 2011.
- L. CONDE, P. M. BRACCI, E. HALPERIN, AND C. F. SKIBOLA. *A Search for Overlapping Genetic Susceptibility Loci Between Non-Hodgkin Lymphoma and Autoimmune Diseases*. Genomics, Vol. 98, No. 1, pp. 9-14, July 2011.
- H. COOK, E. GONINA, S. KAMIL, G. FRIEDLAND, D. PATTERSON, AND A. FOX. *CUDA-Level Performance with Python-Level Productivity for Gaussian Mixture Model Applications*. Proceedings of the Third USENIX Workshop on Hot Topics in Parallelism (HotPar '11), Berkeley, California, May 2011.
- J. FELDMAN AND S. NARAYANAN. *Simulation Semantics, Embodied Construction Grammar, and the Language of Events*. Proceedings of the International Workshop on Language - Action Tools for Cognitive Artificial Agents: Integrating Vision, Action and Language at the 25th Conference on Artificial Intelligence (AAAI-2011), San Francisco, California, August 2011.
- A. GHODSI, M. ZAHARIA, B. HINDMAN, A. KONWINSKI, S. SHENKER, AND I. STOICA. *Dominant Resource Fairness: Fair Allocation of Multiple Resources in Datacenters*. Proceedings of the 8th USENIX Symposium on Networked Systems Design and Implementation (NSDI 2011), Boston, Massachusetts, March 2011.
- M. GRACIARENA, M. DELPLANCHE, E. SHRIBERG, AND A. STOLCKE. *Bird Species Recognition Combining Acoustic and Sequence Modeling*. Proceedings of the IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP 2011), Prague, Czech Republic, pp. 341-344, May 2011.
- Y. HASEGAWA, R. LEE-GOLDMAN, A. KONG, AND K. AKITA. *FrameNet as a Resource for Paraphrase Research*. Constructions and Frames, Special Issue on Advances in Frame Semantics, Vol. 3, Issue 1, pp. 104-127, August 2011.
- D. HERMELIN, M. MNICH, E. J. VAN LEEUWEN, AND G. J. WOEGINGER. *Domination When the Stars Are Out*. Proceedings of the 38th International Colloquium on Automata, Languages and Programming (ICALP 2011), Zurich, Switzerland, pp. 432-473, July 2011.
- B. HINDMAN, A. KONWINSKI, M. ZAHARIA, A. GHODSI, A. JOSEPH, R. KATZ, S. SHENKER, AND I. STOICA. *Mesos: A Platform for Fine-Grained Resource Sharing in the Data Center*. Proceedings of the 8th USENIX Symposium on Networked Systems Design and Implementation (NSDI 2011), Boston, Massachusetts, March 2011.
- L. HODGKINSON AND R. M. KARP. *Algorithms to Detect Multiprotein Modularity Conserved During Evolution*. Proceedings of the 7th International Symposium on Bioinformatics Research and Applications (ISBRA 2011), Changsha, China, May 2011.
- J. HONG AND C. F. BAKER. *How Good Is the Crowd at "Real" WSD? Proceedings of the Fifth Linguistic Annotation Workshop (LAW-V)*, Portland, Oregon, June 2011.
- H. HUNG, G. FRIEDLAND, AND C. YEO. *Computationally Efficient Clustering of Audio-Visual Meeting Data*. In Multimedia Interaction and Intelligent User Interfaces: Principles, Methods, and Applications, M. Etho, J. Luo, and L. Shao, eds., pp. 25-59, June 2011.
- H. HUNG, Y. HUANG, G. FRIEDLAND, AND D. GATICA-PEREZ. *Estimating Dominance in Multi-Party Meetings Using Speaker Diarization from a Single Microphone*. IEEE Transactions on Audio, Speech and Language Processing, Vol. 19, No. 4, pp. 847-860, May 2011.
- C. KANICH, N. WEAVER, D. MCCOY, T. HALVORSON, C. KREIBICH, K. LEVCHENKO, V. PAXSON, G. VOELKER, AND S. SAVAGE. *Show Me the Money: Characterizing Spam-Advertised Revenue*. Proceedings of the 20th USENIX Security Symposium (Security '11), San Francisco, California, August 2011.
- S. KARAYEV, M. FRITZ, S. FIDLER, AND T. DARRELL. *A Probabilistic Model for Recursive Factorized Image Features*. Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR 2011), Colorado Springs, Colorado, pp. 401-408, June 2011.
- R. M. KARP. *Heuristic Algorithms in Computations Molecular Biology*. Journal of Computer and System Sciences, Special Issue: Celebrating Karp's Kyoto Prize, Vol. 77, No. 1, pp. 122-128, January 2011.
- R. M. KARP. *Effective Heuristics for NP-Hard Problems*. Presented at the Michael Rabin Celebration, Cambridge, Massachusetts, August 2011.

# publications, cont.

- B. KINGSBURY, H. SOLTU, G. SAON, S. CHU, H.-K. KUO, L. MANGU, S. RAVURI, A. JANIN, AND N. MORGAN. *The IBM 2009 GALE Arabic Speech Transcription System*. Proceedings of the IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP 2011), Prague, Czech Republic, pp. 4672-4675, May 2011.
- B. KIRKPATRICK, S. C. LI, R. M. KARP, AND E. HALPERIN. *Pedigree Reconstruction Using Identity by Descent*. Proceedings of the 15th Annual International Conference on Research in Computational Molecular Biology (RECOMB 2011), Vancouver, British Columbia, Canada, pp. 136-152, March 2011.
- T. KOPONEN, S. SHENKER, H. BALAKRISHNAN, N. FEAMSTER, I. GANICHEV, A. GHODSI, B. GODFREY, N. McKEOWN, G. PARULKAR, B. RAGHAVAN, J. REXFORD, S. ARIANFAR, AND D. KUPISOV. *Editorial: Architecting for Innovation*. ACM SIGCOMM Computer Communications Review, Vol. 41, No. 3, July 2011.
- O. KRAMER AND F. GIESEKE. *Short-Term Wind Energy Forecasting Using Support Vector Regression*. Proceedings of the International Conference on Soft Computing Models in Industrial and Environmental Applications (SoCo), Salamanca, Spain, April 2011.
- C. KREIBICH, N. WEAVER, C. KANICH, W. CUI, AND V. PAXSON. *GQ: Practical Containment for Measuring Modern Malware Systems*. ICSI Technical Report TR-11-002, May 2011.
- C. KREIBICH, N. WEAVER, G. MAIER, B. NECHAEV, AND V. PAXSON. *Experiences from Netalyzr with Engaging Users in End-System Measurement*. Proceedings of ACM SIGCOMM Workshop on Measurements Up the Stack (W-MUST), Toronto, Canada, pp. 25-30, August 2011.
- B. KULIS, K. SAENKO, AND T. DARRELL. *What You Saw is Not What You Get: Domain Adaptation Using Asymmetric Kernel Transforms*. Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR 2011), Colorado Springs, Colorado, pp. 1785-1792, June 2011.
- V. KWOK, Z. NIU, P. KAY, K. ZHOU, L. MO, Z. JIN, K.-F. SO, AND L. H. TAN. *Learning New Color Names Produces Rapid Increase in Gray Matter in the Intact Adult Human Cortex*. Proceedings of the National Academy of Sciences of the United States of America, Vol. 108, No. 16, pp. 6686-6688, April 2011.
- J. LÄSSIG, B. SATZGER, AND O. KRAMER. *Self-Stabilization in Hierarchically Structured Energy Markets*. Proceedings of the 8th International Conference on Information Technology: New Generations (ITNG 2011), Las Vegas, Nevada, pp. 803-809, April 2011.
- M. LARSON, M. SOLEYMANI, P. SERDYUKOV, S. RUDINAC, C. WARTENA, V. MURDOCK, G. FRIEDLAND, R. ORDELMAN, AND G. J. F. JONES. *Automatic Tagging and Geo-Tagging in Video Collections and Communities*. Proceedings of the ACM International Conference on Multimedia Retrieval (ICMR 2011), Trento, Italy, April 2011.
- Y. LEE, R. AVIZIENIS, A. BISHARA, R. XIA, D. LOCKHART, C. BATTEN, AND K. ASANOVIC. *Exploring the Tradeoffs Between Programmability and Efficiency in Data-Parallel Accelerators*. Proceedings of the International Symposium on Computer Architecture (ISCA-2011), San Jose, California, pp. 129-140, June 2011.
- H. LEI, J. CHOI, AND G. FRIEDLAND. *City-Identification on Flickr Videos Using Acoustic Features*. ICSI Technical Report TR-11-001, Berkeley, California, April 2011.
- H. LEI, J. CHOI, A. JANIN, AND G. FRIEDLAND. *User Verification: Matching the Uploaders of Videos Across Accounts*. Proceedings of the IEEE International Conference on Acoustic, Speech, and Signal Processing (ICASSP 2011), Prague, Czech Republic, pp. 2404-2407, May 2011.
- H. LEI AND N. MIRGHAFORI. *Data Selection with Kurtosis and Nasality Features for Speaker Recognition*. Proceedings of the 12th Annual Conference of the International Speech Communication Association (Interspeech 2011), Florence, Italy, pp. 2753-2756, August 2011.
- K. LEVCHENKO, A. PITSILLIDIS, N. CHACHRA, B. ENRIGHT, M. FELEGYHAZI, C. GRIER, T. HALVORSON, C. KANICH, C. KREIBICH, H. LIU, D. MCCOY, N. WEAVER, V. PAXSON, G. M. VOELKER, AND S. SAVAGE. *Click Trajectories: End-to-End Analysis of the Spam Value Chain*. Proceedings of the 32nd IEEE Symposium on Security and Privacy, Oakland, California, pp. 431-446, May 2011.
- Z. LI, A. GOYAL, Y. CHEN, AND V. PAXSON. *Towards Situational Awareness of Large-Scale Botnet Probing Events*. IEEE Transactions on Information Forensics and Security, Vol. 6, No. 1, pp. 175-188, March 2011.
- D. LOKSHITANOV, M. MNICH, AND S. SAURABH. *Planar k-Path in Subexponential Time and Polynomial Space*. Proceedings of the 37th International Workshop on Graph-Theoretic Concepts in Computer Science (WG '11), Teplá-Kláster, Czech Republic, June 2011.
- G. MAIER, A. FELDMANN, V. PAXSON, R. SOMMER, AND M. VALLENTIN. *An Assessment of Overt Malicious Activity Manifest in Residential Networks*. Proceedings of the 8th Conference on Detection of Intrusions and Malware and Vulnerability Assessment (DIMVA 2011), Amsterdam, The Netherlands, pp. 144-163, July 2011.
- B. T. MEYER, S. V. RAVURI, M. R. SCHAEGLER, AND N. MORGAN. *Comparing Different Flavors of Spectro-Temporal Features for ASR*. Proceedings of the 12th Annual Conference of the International Speech Communication Association (Interspeech 2011), Florence, Italy, pp. 1269-1272, August 2011.
- B. MILLER, P. PEARCE, C. GRIER, C. KREIBICH, AND V. PAXSON. *What's Clicking What? Techniques and Innovations of Today's Clickbots*. Proceedings of the 8th Conference on Detection of Intrusions and Malware & Vulnerability Assessment (DIMVA 2011), Amsterdam, The Netherlands, pp. 164-183, July 2011.
- M. MNICH. *Review of F. V. Fomin and D. Kratsch, "Exact Exponential Algorithms."* Operations Research Letters, Vol. 39, No. 3, pp. 229-230, May 2011.

- L. MO, G. XU, P. KAY, AND L.-H. TANE. Electrophysiological Evidence for the Left-Lateralized Effect of Language on Preattentive Categorical Perception of Color. *Proceedings of the National Academy of Sciences of the United States of America*, Vol. 108, No. 34, pp. 14026-14030, August 2011.
- R. MORRISON, E. FRIEDMAN, AND A. LANDSBERG. Combinatorial Games with a Pass: A Dynamical Systems Approach. *Proceedings of the Fourth International Conference on Chaotic Modeling, Simulation, and Applications (CHAOS2011)*, Agios Nikolaos, Greece, May 2011.
- T. OWENS, K. SAENKO, A. CHAKRABARTI, Y. XIONG, T. ZICKLER, AND T. DARRELL. Learning Object Color Models from Multi-View Constraints. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR 2011)*, Colorado Springs, Colorado, pp. 169-176, June 2011.
- T. OWENS, K. SAENKO, A. CHAKRABARTI, Y. XIONG, T. ZICKLER, AND T. DARRELL. The Ratio Method for Multi-View Color Constancy. *E ECS Department, UC Berkeley, Technical Report No. UCB/E ECS-2011-23*, Berkeley, California, April 2011.
- V. PAXSON. Reflections on Measurement Research: Crooked Lines, Straight Lines, and Moneyshots. Keynote talk at the ACM Special Interest Group on Data Communications Conference (SIGCOMM 2011), Toronto, Canada, August 2011.
- V. PAXSON, M. ALLMAN, J. CHU, AND M. SARGENT. Computing TCP's Retransmission Timer. *RFC 6298*, June 2011
- B. RAGHAVAN AND J. MA. Networking in the Long Emergency. *Proceedings of the ACM SIGCOMM Workshop on Green Networking (GreenNet '11)*, Ontario, Canada, pp. 37-42, August 2011
- S. RAVURI. On the Use of Spectro-Temporal Features in Noise-Additive Speech. UC Berkeley Master's thesis, Spring 2011.
- H. RODRIGUES, J. R. SANTOS, Y. TURNER, P. SOARES, AND D. GUEDES. Gatekeeper: Supporting Bandwidth Guarantees for Multi-Tenant Datacenter Networks. *Proceedings of the Third Workshop on I/O Virtualization (WIOV '11)*, Portland, Oregon, June 2011.
- M. H. SANCHEZ, L. FERRER, E. SHRIBERG, AND A. STOLCKE. Constrained Cepstral Speaker Recognition Using Matched UBM and JFA Training. *Proceedings of the 12th Annual Conference of the International Speech Communication Association (Interspeech 2011)*, Florence, Italy, pp. 141-144, August 2011.
- N. SCHEFFER, L. FERRER, M. GRACIARENA, S. KAJAREKAR, E. SHRIBERG, AND A. STOLCKE. The SRI NIST 2010 Speaker Recognition Evaluation System. *Proceedings of the IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP 2011)*, Prague, Czech Republic, pp. 5292-5295, May 2011.
- E. SHRIBERG AND A. STOLCKE. Language-Independent Constrained Cepstral Features for Speaker Recognition. *Proceedings of the IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP 2011)*, Prague, Czech Republic, pp. 5296-5299, May 2011.
- A. SHYR, T. DARRELL, M. JORDAN, AND R. URTASUN. Supervised Hierarchical Pitman-Yor Process for Natural Scene Segmentation. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR 2011)*, Colorado Springs, Colorado, June 2011.
- A. SINGLA, C.-Y. HONG, L. POPA, AND P. BRIGHTEN GODFREY. Jellyfish: Networking Data Centers Randomly. *Proceedings of the Third USENIX Workshop on Hot Topics in Cloud Computing (USENIX HotCloud '11)*, Portland, Oregon, June 2011
- S. STEIDL, K. RIEDHAMMER, T. BOCKLET, F. HÖNIG, AND E. NÖTH. Java Visual Speech Components for Rapid Application Development of GUI Based Speech Processing Applications. *Proceedings of the 12th Annual Conference of the International Speech Communication Association (Interspeech 2011)*, Florence, Italy, pp. 3257-3260, August 2011.
- A. STOLCKE. Making the Most from Multiple Microphones in Meeting Recognition. *Proceedings of the IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP 2011)*, Prague, Czech Republic, pp. 4992-4995, May 2011
- M. E. TAYLOR, B. KULIS, AND F. SHA. Metric Learning for Reinforcement Learning Agents. *Proceedings of the 10th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2011)*, Taipei, Taiwan, May 2011.
- K. THOMAS, C. GRIER, J. MA, V. PAXSON, AND D. SONG. Design and Evaluation of a Real-Time URL Spam Filtering Service. *Proceedings of the 32nd IEEE Symposium on Security and Privacy*, Oakland, California, pp. 447-464, May 2011.
- D. VALLET, F. HOPFGARTNER, J. M. JOSE, AND P. CASTELLS. Effects of Usage Based Feedback on Video Retrieval: A Simulation Based Study. *ACM Transactions on Information Systems*, Vol. 29, Issue 2, Article 11, April 2011.
- O. VINYALS AND S. RAVURI. Comparing Multilayer Perceptron to Deep Belief Network Tandem Features for Robust ASR. *Proceedings of the 36th IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP '11)*, Prague, Czech Republic, pp. 4596-4599, May 2011.
- A. WATERMAN, Y. LEE, D. PATTERSON, AND K. ASANOVIC. The RISC-V Instruction Set Manual, Volume I: Base User-Level ISA. *E ECS Department, UC Berkeley, Technical Report No. UCB/E ECS-2011-62*, Berkeley, California, May 2011.
- N. WEAVER, C. KREIBICH, B. NECHAEV, AND V. PAXSON. Implications of Netalyzr's DNS Measurements. *Proceedings of the First Workshop on Securing and Trusting Internet Names (SATIN)*, Teddington, United Kingdom, April 2011.
- N. WEAVER, C. KREIBICH, AND V. PAXSON. Redirecting DNS for Ads and Profit. *Proceedings of the 20th USENIX Security Symposium's Workshop on Free and Open Communications on the Internet (FOCI '11)*, San Francisco, California, August 2011.
- M. ZAHARIA, M. CHOWDHURY, T. DAS, A. DAVE, J. MA, M. MCCAULEY, M. FRANKLIN, S. SHENKER, AND I. STOICA. Resilient Distributed Datasets: A Fault-Tolerant Abstraction for In-Memory Cluster Computing. *E ECS Department, UC Berkeley, Technical Report No. UCB/E ECS-2011-82*, Berkeley, California, July 2011.
- M. ZAHARIA, B. HINDMAN, A. KONWINSKI, A. GHODSI, A. JOSEPH, R. KATZ, S. SHENKER, AND I. STOICA. The Datacenter Needs an Operating System. *Proceedings of the Third USENIX Workshop on Hot Topics in Cloud Computing (Hot Cloud '11)*, Portland, Oregon, June 2011.



INTERNATIONAL  
COMPUTER SCIENCE  
INSTITUTE

THE INTERNATIONAL COMPUTER SCIENCE INSTITUTE (ICSI) is one of the few independent, non-profit US labs solely conducting open, non-proprietary, pre-competitive research in computer science. ICSI's mission is to further research in computer science through international collaboration, and further international collaboration through research in computer science. Affiliated with the University of California campus in Berkeley, 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 these 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; natural speech processing; computer vision; and computer architecture.

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 used without explicit prior written permission.

BOARD OF TRUSTEES

Hervé Bourlard  
David Culler  
Deborah Estrin  
Graham Fleming  
Clifford Higgerson  
Nelson Morgan  
David Nagel  
Peter Norvig  
Prabhakar Raghavan  
Scott Shenker  
Eero Silvennoinen  
David Tennenhouse  
Wolfgang Wahlster

EDITOR-IN-CHIEF

Maria Quintana

SENIOR EDITOR

Nelson Morgan

COPYWRITER

Chris Switzer

LAYOUT

Jacob Wolkenhauer

PHOTOGRAPHY

Jaci Considine

INTERNATIONAL  
COMPUTER  
SCIENCE  
INSTITUTE

1947 Center Street Suite 600  
Berkeley, CA 94704 USA  
v 510 666 2900  
[www.icsi.berkeley.edu](http://www.icsi.berkeley.edu)