the newsletter of the International Computer Science Institute

The ICSI GAZETTE

volume eleven | issue two | summer 2013

featured researcher: nelson morgan

Morgan has led speech research at ICSI since the Institute's inauguration in 1988. Morgan also served as director for thirteen years starting in 1999, the year the agreement that had established ICSI expired. Morgan volunteered for the challenge of broadening and stabilizing the Institute's funding base and, by the end of his tenure, left the Institute in better financial condition than it had experienced in years. Morgan has always enjoyed a challenge.

Early Career

Morgan was born and raised in Buffalo, New York. Fascinated by electronics, he was "always hooking things up to other things," he said. When he was 10, he built a simple noise recognizer for Halloween by connecting a microphone to a circuit board he had bought from Popular Electronics. When the microphone detected sound, a tape player began playing Godzilla music to scare trick-or-treaters.

Morgan's older brother had a strong influence on his early life and interests: he introduced Morgan to rock-and-roll, which sparked a lifelong interest in sound and audio, and gave him his first tape recorder, a rare technology in the 1950s. Morgan used it to record televised news reports, which he edited – using a razor blade – to make them sound more leftist. Morgan's brother also introduced him to the college programs that allowed him to enter college after two years of high school. At the age of 16, Morgan entered the University of Chicago as a physics major.

But he decided to leave the university in order to do "a lot of wandering," at one point living in a teepee in the woods. He eventually began managing and recording rock bands and working as a technical advisor on motion pictures, including *Godfather Part II*.

Most of his work consisted of mixing audio for television commercials. One day, while adjusting the volume of a dog's bark for a commercial, he said, "I realized this was not quite the creative experience I was expecting."

The Science of Sound

He thought he might enjoy the audio-related career he had wandered into if he better understood the science behind it, so he began taking classes part time while continuing to work as a sound technician. One summer, he took an introductory course in electronics at UC Berkeley. "It was wonderful," he said. "It was much more exciting than anything I was doing in the studio."

His undergraduate advisor at Berkeley suggested he apply for a National Science Foundation fellowship, which would allow him to attend school full time as a graduate student. "I decided I would write up exactly what I wanted to do – just what I wanted to do – for my research," he said. "And if they said they'd pay for it, great, I'd be a student. And if not, I would continue doing what I was doing."

His fellowship application described a project to create

continued on page 4

in this issue:

photo by Leide Porcu

reflections by Roberto Pieraccini, CEO

Of the many words that characterize the last few years of technical innovation, three stand out: "big," as in big data; "many," as in the use of many people to help make sense of big data; and "deep," as in deep learning. What is big about big data is the notion that, in an unprecedented way, we have access to amounts of data so large that we are facing not only new and seemingly insurmountable problems, but also great, or rather big, new opportunities for a deeper understanding of the phenomena represented by data. The problems derive from the sheer size of the ever-growing available data: simply storing, handling, and analyzing massive amounts of data is a problem in itself. The opportunities, however, are endless, and they promise to be epochal. If we can process and extract the appropriate knowledge from these large amounts of data, we can answer the most important unanswered questions about our world and create new tools that would benefit society at large.

Just look, for instance, at the growth rate of the data that we most frequently access. The numbers are mindboggling: as just one example, YouTube recently reported that 72 hours of video are uploaded to its Web site every minute! That's big data. Finding a video, unless it is properly tagged with text, is like finding a needle in a haystack. While searching huge amounts of text can be done in real time — and we do this every time we run a Web search — searching for visual concepts within an enormous and ever-growing corpus of video and audio may be beyond our current computational capability. Finding out how to do this with what we have is a difficult problem to solve, but the solution would give us the unprecedented opportunity to take advantage of all the data, video, and audio - and not just text — that is available on the Web for deeper understanding and deeper searches.

In order for machines to learn how to use big data, we need some level of supervision from humans, known as the annotation process in machine learning. Most machine learning techniques start with a sizable amount of human-annotated data, in which every sample is associated with some form of truth or knowledge representation. Often many non-experts are able to annotate certain types of data, like video or audio. With the advent of crowd-sourcing, for instance with Amazon's Mechanical Turk, it is possible to automate the whole process: crowd enrollment, data distribution, payment, and management. Machine learning complements the annotation work of crowds and increases its efficiency and accuracy. Data alone, even big data, is not enough for many tasks. The possibility of having big data annotated with the help of machines and crowds in order to come to a deeper understanding of our world is huge. That's the effort undertaken by ambitious initiatives such as UC Berkeley's AMPlab, which includes some of ICSI's affiliated researchers, and by other research initiatives at ICSI.

Genomics is another area that benefits from the availability of big data. The cost of DNA sequencing is predicted to decline at a rate of half every six months, and even though that rate can slow temporarily because of economic forces, it is still a much faster pace than the growth of computational power predicted by Moore's law. If that holds, the amount of genomic data available either from a single individual or from an entire population is bound to grow disproportionately compared with our ability to process it in a timely manner. That's the problem, and unless we find ways to speed up the computing power of our machines much faster than Moore's law, we need to work hard at trying to understand new ways, new algorithms, to manipulate that vast amount of data. But the potential rewards are enormous, including that of finding personalized treatments for cancer that are more effective and more timely than they are presently and that can help increase the chances of patient survival. In other words, we must use BIG data for a DEEP understanding of our world, and MANY are helping toward that goal.

new research group: audio and multimedia

ICSI has established a new research area, Audio and Multimedia, to focus on problems related to multimedia analysis and retrieval. Researchers in this area, led by Gerald Friedland, are working on ways to extract meaning automatically from the vast amounts of consumer-produced data available freely on Web sites like Flickr and YouTube – a dataset of billions of images and videos with few constraints on quality, size, and content. Videos are of particular interest to the group because they provide textual, audio, and visual information for analysis and are the fastest-growing type of content on the Internet. YouTube claims that 72 hours of video are uploaded to its Web site alone every minute.

Gerald Friedland, Audio and Multimedia Director

Consumer-produced content consists of entertainment, instructions, personal records, and aspects of life as it was when the media was recorded. They represent a compendium of information about trends, phenomena or events, and social context and dynamics, and so are useful for qualitative and quantitative empirical research on a larger scale than has ever been possible before. The hope is that multimedia research will enable tools to easily organize and search these large collections of social media content and to aid this research in both academic and industrial contexts.

Evolved from the Speech Group, Audio and Multimedia research puts a special focus on audio analysis. Audio content is frequently complementary to visual content, as in videos, but has received less attention from the multimedia research community.

For example, in the ALADDIN project, funded by IARPA, researchers are building a system that uses acoustic analysis to search for concepts in videos. IARPA has provided ICSI, along with other teams from institutions around the world, 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 150,000 unlabeled videos. The team is developing parallelization methods to deal with the immense data and was recently given multi-CPU cards by Intel for this purpose. ICSI's work is in collaboration with SRI, Carnegie Mellon University, and other research institutions.

A major focus of Audio and Multimedia researchers has been finding the origin of videos and this led to the development of a method for qualifying Mechanical Turk users. Mechanical Turk is Amazon's answer to crowd-sourcing, in which average users complete a task for small amounts of money. Crowd-sourcing is used most often to complete tasks that are easy for humans and difficult for computers, but location estimation is difficult for both. The researchers wanted to be able to find a group of people who were qualified to estimate the locations of videos, partly to give a human baseline to compare their system to. They were also developing methods to qualify Mechanical Turk participants so that they could eventually conduct a Mechanical Turk task to place videos from areas that are not well represented on the Web, such as West Africa.

To identify qualified people, they asked participants to estimate the location of a series of videos using online resources such as Google Maps and image search. Their results were compared to those achieved by researchers in the group and by researchers in the multimedia community. They found that about a fifth of their participants were qualified to place videos, meaning that they were able to place videos within 10 kilometers of their origins 80 percent of the time.

The ability to automatically find the origins of videos has implications for online privacy, and the researchers of Audio and Multimedia, along with Robin Sommer and Nick Weaver of Networking and Security, are working to expose the dangers. One project focuses on teaching high school students about the privacy implications of sharing content online since younger people, as the group that uses social media the most and is also the least aware of its potential consequences, are particularly vulnerable to attacks. The researchers are developing classroom tools and building a Web site that help

continued on page 9

featured researcher, continued

continued from page 1

sound effects electronically. At the time, technicians commonly simulated room reverberation with a metal plate in order to produce sound effects for movies. "These were pretty hokey-sounding, and you couldn't adjust them for a particular room size," he said.

He was awarded a three-year graduate fellowship from NSF, and he decided he would try to finish his doctorate before the funding ran out. Doctorates in physical sciences or engineering often take five years or more, but Morgan said he had an advantage: "From the first day I knew exactly what my research would be."

Although his research was on room acoustics, he spent some of his spare time chatting about technical topics with Ben Gold, a pioneer in digital signal processing and then a visiting professor at UC Berkeley. Later, in the early 1990s, Gold and Morgan established a class at Berkeley that combined their varied experiences with speech processing. The class has been taught every other year since then, and Morgan and Gold developed the class outlines into a textbook, *Speech and Audio Signal Processing*, which was recently revised and released in a second edition with the help of Columbia Professor and ICSI alum Dan Ellis.

Neural Networks

As he approached graduation, Morgan was offered a position by Dolby to start a digital audio laboratory. The recession of the late 1970s, however, forced Dolby to lay off much of its work force, and Morgan's offer was canceled at the last moment. He quickly found a position at National Semiconductor, where he worked on speech analysis and synthesis techniques.

In one project at the lab, short recordings of actors were used to synthesize longer pieces of audio for commercials and products such as talking soda machines. This required that the recordings be divided into voiced speech, which is produced by vibrations of the vocal cords, and unvoiced speech, which is produced from air moving past some obstruction in the vocal tract such as the teeth. To explore methods of separating these automatically, Morgan bought a book on pattern recognition and coded all the techniques in it. Neural networks, which Morgan would later use extensively in speech recognition, happened to work the best in his experiments. "We cut the time enormously by just having experts do fine-tuning on a smaller set and training the classifier from the hand-labeled data," he said.

His next experience with neural networks was at EEG Labs, which he joined in 1984. Researchers at the lab were using scans of the brain in order to understand its performance of cognitive functions. It was a new experience for Morgan, whose work until then had been mostly in signal processing. "I learned a lot from them, not just about the brain, but also about pattern recognition and neural networks," Morgan said. "That's really where I learned about them."

In 1986, a new computer science laboratory was incorporated in Berkeley, and word got around that it needed researchers. After a conversation with then-Director Jerry Feldman, Morgan was chosen to lead a group at ICSI that would focus on building massively parallel computers.

"But I didn't want to be just building up systems to do what someone else was interested in," he said. He decided the group's work would be applied to problems in speech research. In September 1988, when the Institute was officially inaugurated, he became the leader of the Realization Group, renamed the Speech Group in 1999.

The Realization Group

The group's early successes were in designing and building machines powerful enough to do speech recognition. In 1989, the group designed an array of digital signal processing chips in a ring topology that used programmable gate arrays to interconnect processors. The Ring Array Processor (RAP) had a simple architecture and could be built from off-the-shelf materials. "It was way faster for what we were doing than anything you could buy for any reasonable amount of money," said Morgan.

The RAP, as well as other hardware designed by the group, was used by ICSI's research partners around the world. Sharing of hardware is common now, with inexpensive and universal components readily available, but it was unusual in the late 1980s and 1990s.

The group designed computer architectures (including the first single chip vector microprocessor, designed by then-student Krste Asanovic), and built hardware and software. Still, Morgan said, "We had this undercurrent of speech work as being the end goal."

Front-End Processing

By the early 1990s, Morgan's work focused on speech recognition algorithms, rather than on the devices to implement the algorithms. While his efforts at ICSI began with neural network approaches to speech recognition, he began to also work seriously on front-end speech processing: processing of the audio features that are fed to the statistical engine. "I became a real advocate of the idea that you should pay a lot more attention to the front end than automatic speech recognition researchers usually do," he said.

In 1990, the Institute hosted the Speech Recognition Front End Workshop. There, Jordan Cohen, who would later become a frequent ICSI collaborator, presented the "Problem of the Inverse E": if you build a system to filter out the spectrum of the sound "E" from a speech data set, a human listener can still hear the "E's." Morgan realized that human hearing must be sensitive to the transition between sounds so that fixed spectral changes might not eliminate the perception of speech categories. He and his colleague Hynek Hermansky figured that speech recognition systems would do well to process features relatively. They eventually developed this idea into the relative spectral processing technique (RASTA). This kind of processing helps machines handle changes in the audio spectrum. At the time, most speech systems had difficulty, for example, dealing with audio recorded on different microphones from those used to record its training data. This became particularly important later, when cell phones were ubiquitous - RASTA was designed into millions of phones.

This technique and other algorithmic developments at ICSI were used in ICSI's Berkeley Restaurant Project, a spoken dialog system that gave restaurant recommendations. The system was unusual in that both the system and its users could direct the next step in the dialog, and the system could continue a conversation even when users did not respond directly to its questions.

Importantly, said Morgan, the work on RASTA features, as well as more recent successes, stressed the importance of front-end processing. "We woke people up to the fact that training-test spectral mismatch was a problem," he said. "We weren't the first people to suggest that, but we may have been the first ones to talk about it so loudly."

RASTA is also an example of technology emulating human systems, a theme throughout much of Morgan's work. "It's really important to pay attention to what mechanisms we can discover from biological systems," he said.

Starting in 1988, Morgan also collaborated with Hervé Bourlard, the Institute's first visiting scholar, on developing the hybrid approach to speech processing. In this approach, the acoustic probabilities of hidden Markov models (HMMs), which have long been used in speech recognition, are determined through artificial neural networks, which comprise nodes that communicate through connectionist models. Bourlard and Morgan's paper on the approach won a best paper award from the *IEEE Signal Processing Magazine* in 1996, and their work together inspired other research directions throughout the 1990s. The hybrid approach is experiencing a comeback with the growing popularity of work on deep learning.

"Working with Morgan is always fun. When you come up with a new idea, he often disagrees and argues with you," said Bourlard, who now sits on ICSI's board of trustees and leads Idiap in Martigny, Switzerland. "That's when you know that you may have got something interesting, and that there may also be more to it."

ICSI Meeting Corpus

By the late 1990s, the Speech Group was looking for more difficult problems. Morgan said, "We were mostly looking at robustness in some sense – why are speech recognition systems breaking down? How do you make them less sensitive?"

A student suggested that Morgan, who was on his way to a meeting in Europe, keep notes about when a handheld speech recognition system, such as Siri, would have been useful. Morgan realized he needed, not a personal electronic assistant,

continued on next

continued from previous

but some easy way of recording and retrieving notes from the meeting.

"All the sudden it struck me: that's the key application. You want to be able to have access to information from some extended meeting or meetings by querying for it," he said.

From this idea emerged the ICSI Meeting Corpus, a collection of recorded audio from meetings held at the Institute, along with transcriptions to aid in training speech recognition systems. At the time, it was the largest corpus of publicly available transcribed meetings.

It was important that these recordings were of spontaneous speech. They included laughter, speech from multiple people talking at the same time, and vocalized pauses – "ums" and so forth. These elements, said Morgan, made for interesting problems in speech recognition, which the team set about solving.

The Next Challenge

While the Speech Group was looking for challenges in the late 1990s, the Institute had its own. Jerry Feldman, the Institute's first director, was planning to step down from the position and, at around the same time, the funding agreement with Germany that had established the Institute in 1986 was about to run out. There were discussions about whether the Institute would close its doors.

"That just seemed like such a waste to me," Morgan said. "It just felt like there was so much here that was good."

Morgan volunteered to take over directorship of the Institute, but the financial situation was grim. "Morgan did not take this job out of ambition, but out of duty," said Scott Shenker, director of Research Initiatives and ICSI's chief scientist.

With the reduction in international funding, the Institute had to find industrial and U.S. federal support. A major source of revenue was AT&T, which funded a new center at ICSI focused on Internet research. Shenker helped draw the center to the Institute.

Over the next few years, ICSI had a balanced budget sheet, with funding from industrial, U.S. federal, and some international partners. The outlook got even brighter when Richard Karp, formerly the Algorithms Group leader and a Turing Award winner, returned from a four-year visit to the University of Washington.

But the dot-com bust of the early 2000s led to significant reductions in industrial funding. Since then, the Institute has come to rely mainly on federal support, particularly from the National Science Foundation and other Federal sources. Additional support comes from industry and international partners.

A major accomplishment was the establishment of a new German visiting agreement. The Institute's original ten-year agreement with the German Federal Laboratory for Computer Science expired in 1999. Morgan negotiated new agreements with Germany organized through the German Academic Exchange Service. In its most recent form, the agreement supports the hosting of about ten postdoctoral fellows every year from Germany. The Institute currently also has agreements with Finland and Singapore. Morgan was instrumental in all of them. Under Morgan's leadership, the Institute also recently received several large federal U.S. grants.

> "Morgan's first love is research, but he sacrificed the pursuit of his own intellectual agenda in order to provide financial stability for the rest of us."
>
> - Scott Shenker

Since 1992, Morgan has also held a faculty position in the Electrical Engineering and Computer Science Department. He has advised 20 doctoral students in that time.

Last year, Morgan stepped down from the position as director and now serves as deputy director. He will continue, as he has done since the Institute's foundation, to work on topics in speech, and his focus is gradually switching back to research.

"Morgan's first love is research, but he sacrificed the pursuit of his own intellectual agenda in order to provide financial stability for the rest of us," said Shenker. "He did so with quiet grace and relentless energy, and we are all in his debt."

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

AI
Gerard de Melo (Germany)
Aida Nematzadeh

Algorithms
Michael Elberfeld (Germany)
Itamar Eskin
Farhad Hormozdiari

Architecture
Miquel Moretó Planas

Speech
Hai Do (Singapore)
Arlo Faria (Brazil)
Daniel Göhring (Germany)
Frantisek Grezl
David Imseng (Switzerland)
Kalle Palomaeki (Finland)
Mirco Ravanelli
Korbinian Riedhammer (Germany)

Networking
Bernhard Amann (Germany)
Jason Croft
Haixin Duan
Andrei Gurtov (Finland)
Shuang Hao
Jussi Kangasharju (Finland)
Sheharbano Khattak
Dmitriy Kuptsov (Finland)
Christof Leng (Germany)
Andrey Lukyanenko (Finland)
Ilya Nikolaevsky (Finland)
André Schumacher (Finland)
Amin Tootoonchian
Daniel Warneke (Germany)
Andreas Wundsam (Germany)
Jiao Zhang

Campus Affiliation / Other
Nils Peters (Germany)

Vision
Stefanie Jegelka
Matthias Kirchner (Germany)
Lorenzo Riano
Erik Rodner (Germany)


Jussi Kangasharju

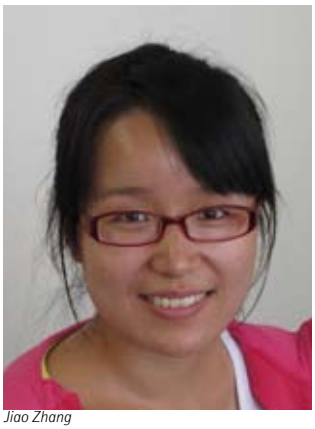
Jiao Zhang


Aida Nematzadeh


Christof Leng

news briefs

Algorithms Leader Richard Karp Moderates Panel on Alan Turing's Work

September 5, 2012

Richard Karp, former Algorithms leader, moderated a panel discussion on Alan Turing's influence on current research in logic, computer science, complexity, and biology on September 4 in the Berkeley City College Auditorium. Karp received the Turing Award in 1985 for his contributions to the theory of computational algorithms. The event was part of a year-long celebration of Turing's 100th birthday and was hosted by MSRI.

Vern Paxson Is New Director of Networking and Security; Steven Wegmann Is New Director of Speech

January 28, 2013

Professor Vern Paxson is the director of Networking and Security research. Paxson was previously a senior researcher in Networking. He takes the role of director over from Professor Scott Shenker, who directs the recently established Research Initiatives area and serves as ICSI's chief scientist. In addition, Steven Wegmann is the new director of Speech research. Professor Nelson Morgan, the founder and leader of the group for more than 24 years, will be focusing on research activities, passing on the leadership of the team to Wegmann.

Speech Researcher Nikki Mirghafori Selected to Attend NAKFI Conference

November 20, 2012

Nikki Mirghafori, a senior Speech researcher, attended the prestigious National Academies Keck Futures Initiative (NAKFI) conference November 14-17 in Irvine, California. Participants in the conference, which is held once a year, are selected through a highly competitive process and come from a range of scientific backgrounds.

Board Member Receives IJCAI Donald E. Walker Distinguished Service Award

February 20, 2013

Professor Wolfgang Wahlster, a member of the ICSI Board of Trustees, is the recipient of the 2013 IJCAI Donald E. Walker Distinguished Service Award for his substantial contributions and his extensive service to the field of artificial intelligence throughout his career. The award was presented during the opening ceremony of the 23rd International Joint Conference on ArtificialIntelligence in Beijing, China in August. Wahlster is the director and CEO of the German Research Center for Artificial Intelligence (DFKI GmbH) and a professor of computer science at Saarland University.

New Service Checks Users' SSL Certificates

November 7, 2012

Networking researchers have made available a new service that provides near real-time reputation information on TLS/SSL certificates. The ICSI Certificate Notary improves the security of users' Internet activity by allowing clients to compare certificates they encounter while browsing the Web against a third-party database. Such checks help avoid man-in-the-middle attacks, where an adversary assumes the identity of a well-known secure site (for example, a bank or email provider) by serving a malicious certificate. SlashDot featured an interactive graph created by DAAD postdoctoral fellow Bernhard Amann based on data collected for the notary on December 14.

ICSI Appoints New Members to Board of Trustees

October 15, 2012

ICSI has announced the appointment of six members to its Board of Trustees, growing the diverse group of academic and corporate leaders to 16 members. The expanded board held its annual meeting on October 12, when the new trustees officially assumed their duties. The new members are Greg Badros of Facebook; Deborah Crawford of Drexel University; Mazin Gilbert of AT&T Labs; Barbara Grosz of Harvard University; Eric Horvitz of Microsoft Research; and Martti Mäntylä of EIT ICT Labs.

Networking Researcher Sylvia Ratnasamy Receives Yahoo! ACE Award

November 26, 2012

Sylvia Ratnasamy has received a Yahoo! Academic Career Enhancement (ACE) Award for the 2012-2013 academic year. Ratnasamy is a researcher in Networking and Security and an assistant professor at UC Berkeley. Her research focuses on the design and implementation of networked systems.

ICSI Team Builds Speech Recognition Systems for New Languages Under Severe Constraints

November 28, 2012

ICSI is leading a research team under the IARPA Babel Program that is building speech recognition solutions with self-imposed time and data limitations for a variety of languages. The work aims to better understand fundamental challenges and discover new methods of developing speech models for languages that could emerge as important in the future.

Video Location Estimation System Receives Honorable Mention at MediaEval

October 9, 2012

Speech researchers and their collaborators at UC Berkeley received distinctive mention at the MediaEval 2012 Workshop for their video location estimation system. Their system was named the most novel theoretical approach to addressing the workshop's placing task, which challenged participants to find the latitude and longitude of consumer-produced videos and photos pulled from Flickr and YouTube.

icsi in the press

Algorithms developed by Audio and Multimedia researchers are able to estimate where videos were taken.

- "Nowhere to Hide: Video Location Tech Has Arrived," Feb. 21, 2013, Tech's Bottom Line, InfoWorld, Bill Snyder.
- "Researchers Find Way to Pinpoint Where Online Video Was Shot," Feb. 19, 2013, Consumer Tech Radar, CIO Blogs, Bill Snyder.

An interactive graph created by DAAD postdoctoral fellow Bernhard Amann shows the relationships between root and intermediate certificate authorities of institutions around the world.

- "An Interactive Graph of the Certificate Authority Ecosystem," December 14, 2012, Slashdot.

A crackdown by Visa on affiliate programs that sell pharmaceuticals and counterfeit goods over the Internet has been effective in reducing their profits.

- "Rogue Pharma, Fake AV Vendors Feel Credit Card Crunch," October 18, 2012, Krebs on Security, Brian Krebs.

Deputy Director and Speech researcher Nelson Morgan comments on recent advances in speech technology and the future of speech research.

- "Promises and Promises," October 11, 2012, Nelson Morgan, contributed article, Speech Technology.

NSF awards $10 million grant to ICSI, UC San Diego, and George Mason University to study the role played by economics and social interactions in cybercrime.

- "Colleges Receive $10 Million Grant to Study Cyber Crime," September 28, 2012, U.S. News University Connection, Catherine Groux.
- "Cybercrime Project Receives $10 Million from NSF," September 28, 2012, Daily Cal, Caroline Murphy.
- "Exploring Cybercriminal Minds, Safeguarding Privacy Among $50M Worth of New NSF Research Projects," September 27, 2012, Network World, Bob Brown.
- "NSF Grant to Help Understand the World of Cybercrime," September 27, 2012, DFI News.
- "UC Berkeley Part of $10 Million NSF Study of 'Human Element' in Cybercrime," September 25, 2012, San Francisco Business Times, Steven Brown.
- "Berkeley Researchers May Track Twitter Users," September 25, 2012, EETimes, Rick Merritt.
- "Grant to Help Computer Scientists Understand the World of Cybercrime," September 25, 2012, Phys.org.

Audio and Multimedia Director Gerald Friedland comments on publicly available data stockpiled by government.

- "Free Public Database Access Encourages New Applications, Abuses," November 21, 2012, CBS5, Alan Martin.

publications

M. Allman. Comments on Bufferbloat. ACM SIGCOMM Computer Communications, Vol. 3, No. 1, pp. 30-37, January 2013.

B. Amann, R. Sommer, A. Sharma, and S. Hall. A Lone Wolf No More: Supporting Network Intrusion Detection with Real-Time Intelligence. Proceedings of the 15th International Symposium on Attacks, Intrusions, and Detections (RAID 2012), pp. 314-333, Amsterdam, the Netherlands, September 2012.

B. Amann, M. Vallentin, S. Hall, and R. Sommer. Extracting Certificates from Live Traffic: A Near Real Time SSL Notary Service. ICSI Technical Report TR-12-014, November 2012.

B. Amann, M. Vallentin, S. Hall, and R. Sommer. Revisiting SSL: A Large Scale Study of the Internet's Most Trusted Protocol. ICSI Technical Report TR-12-015, December 2012.

S. Beamer, K. Asanovic, and D. Patterson. Direction-Optimizing Breadth-First Search. Proceedings of the International Conference for High Performance Computing, Networking, Storage, and Analysis (SC12), Article No. 12, Salt Lake City, Utah, November 2012.

R. Beverly and M. Allman. Findings and Implications from Data Mining the IMC Review Process. ACM SIGCOMM Computer Communication Review, Vol. 3, No. 1, pp. 23-29, January 2013.

C. Boehm, G. de Melo, F. Naumann, and G. Weikum. LINDA: Distributed Web-of-Data-Scale Entity Matching. Proceedings of the 21st ACM International Conference on Information and Knowledge Management (CIKM 2012), Maui, Hawaii, December 2012.

J. Choi, V. Ekambaram, G. Friedland, and K. Ramchandran. The 2012 ICSI/Berkeley Video Location Estimation System. Presented at the MediaEval 2012 Workshop, Pisa, Italy, October 2012.

M. Dhawan, J. Samuel, R. Teixeira, C. Kreibich, M. Allman, N. Weaver, and V. Paxson. Fathom: A Browser-based Network Measurement Platform. Proceedings of the 2012 Internet Measurement Conference (IMC 2012), Boston, Massachusetts, November 2012.

B. Elizalde and G. Friedland. Taming the Wild: Acoustic Segmentation in Consumer-Produced Videos. ICSI Technical Report TR-12-016, January 2013.

B. Elizalde, G. Friedland, H. Lei, and A. Divakaran. There is No Data Like Less Data: Percepts for Video Concept Detection on Consumer-Produced Media. Proceedings of the ACM International Workshop on Audio and Multimedia Methods for Large-Scale Video Analysis (AMVA) at ACM Multimedia 2012 (MM'12), Nara, Japan, pp. 27-32, October 2012.

J. Feldman. The Neural Binding Problem(s). Cognitive Neurodynamics, Vol. 7, Issue 1, pp. 1-11, February 2013.

C. J. Fillmore. Encounters with Language. Computational Linguistics, Vol. 38, No. 4, pg.701-718, December 2012.

A. Freytag, E. Rodner, P. Bodesheim, and J. Denzler. Beyond Classification -- Large-Scale Gaussian Process Inference and Uncertainty Prediction. Presented at the BigVision Workshop at the 26th Annual Conference on Neural Information Processing Systems (NIPS 2012), Lake Tahoe, Nevada, December 2012.

A. Freytag, E. Rodner, P. Bodesheim, and J. Denzler. Rapid Uncertainty Computation with Gaussian Processes and Histogram Intersection Kernels. Proceedings of the 11th Asian Conference on Computer Vision, Daejeon, Korea. Recipient of a best paper honorable mention award. November 2012.

T. Gloe, S. Pfennig, and M. Kirchner. Unexpected Artefacts in PRNU-Based Camera Identification: A "Dresden Image Database" Case-Study. ACM Multimedia and Security Workshop, Coventry, UK, September 2012.

L. Gottlieb, J. Choi, P. Kelm, T. Sikora, and G. Friedland. Pushing the Limits of Mechanical Turk: Qualifying the Crowd for Video Geo-Location. Proceedings of the ACM Workshop on Crowdsourcing for Multimedia (CrowdMM 2012), held in conjunction with ACM Multimedia 2012, pp. 23-28, Nara, Japan, October 2012.

C. Grier, L. Ballard, J. Caballero, N. Chachra, C. J. Dietrich, K. Levchenko, P. Mavrommatis, D. McCoy, A. Nappa, A. Pitsillidis, N. Provos, M. Zubair Rafique, M. Abu Rajab, C. Rossow, K. Thomas, V. Paxson, S. Savage, and G. M. Voelker. Manufacturing Compromise: The Emergence of Exploit-as-a-Service. Proceedings of the 19th ACM Conference on Computer and Communications Security (CCS 2012), pp. 821-832, Raleigh, North Carolina, October 2012.

M. Gymrek, A. L. McGuire, D. Golan, E. Halperin, and Y. Erlich. Identifying Personal Genomes by Surname Inference. Science, Vol. 339, No. 6117, pp. 321-324, January 2013.

J. Hoffman, K. Saenko, B. Kulis, and T. Darrell. Discovering Latent Domains for Multisource Domain Adaptation. Proceedings of the 12th European Conference on Computer Vision (ECCV 2012), pp. 702-715, Firenze, Italy, October 2012.

A. Janin, G. Friedland, and R. Sommer. Content-Based Privacy for Consumer-Produced Multimedia. ICSI Technical Report TR-12-012, October 2012.

S. Karayev, T. Baumgartner, M. Fritz, and T. Darrell. Timely Object Recognition. Presented at the 26th Annual Conference on Neural Information Processing Systems (NIPS 2012), Lake Tahoe, Nevada, December 2012.

M. Kirchner, P. Winkler, and H. Farid. Impeding Forgers at Photo Inception. Proceedings of SPIE-IS&T Electronic Imaging, Media Watermarking, Security, and Forensics 2013, SPIE Vol. 8665, February 2013.

M. T. Knox, N. Mirghafori, and G. Friedland. Where Did I Go Wrong? Identifying Troublesome Segments for Speaker Diarization Systems. Proceedings of the 13th Annual Conference of the International Speech Communication Association (InterSpeech 2012), Portland, Oregon, September 2012.

J. Makin and S. Narayanan. Real-Time Control of Human Coagulation. IET Control Theory and Applications, Vol. 6, Issue 17, pp. 2630-2643, November 2012.

D. McCoy, H. Dharmdasani, C. Kreibich, G. M. Voelker, and S. Savage. Priceless: The Role of Payments in Abuse-Advertised Goods. Proceedings of the 19th ACM Conference on Computer and Communications Society (CCS 2012), pp. 845-856, Raleigh, North Carolina, October 2012.

G. de Melo. Exploiting Web Data Sources for Advanced NLP. Proceedings of the 24th International Conference on Computational Linguistics (COLING), Tutorial Notes, Mumbai, India, December 2012.

B. Meyer, C. Spille, B. Kollmeier, and N. Morgan. Hooking Up Spectro-Temporal Filters with Auditory-Inspired Representations for Robust Automatic Speech Recognition. Proceedings of the 13th Annual Conference of the International Speech Communication Association (InterSpeech 2012), Portland, Oregon, September 2012.

J. Oliva, J. Feldman, L. Giraldi, and E. Dodge. Ontology Driven Contextual Reference Resolution in Embodied Construction Grammar. Proceedings of the 7th Annual Constraint Solving and Language Processing Workshop (CSLP 2012), Orléans, France, September 2012.

C. de Paiva, A. Rademaker, and G. de Melo. OpenWordNet-PT: An Open Brazilian WordNet for Reasoning. Demo at the 24th International Conference on Computational Linguistics (COLING), Mumbai, India, December 2012.

N. Peters, H. Lei, and J. Choi. Matching Artificial Reverb Settings to Unknown Room Recordings: A Recommendation System for Reverb Plugins. Proceedings of the 133th Audio Engineering Society (AES) Convention, San Francisco, California, October 2012.

N. Peters, H. Lei, and G. Friedland. Name That Room: Room Identification Using Acoustic Features in a Recording. Proceedings of ACM Multimedia 2012, pp. 841-844, Nara, Japan, October 2012.

B. Raghavan and S. Hasan. Macroscopically Sustainable Networking: An Internet Quine. ICSI Technical Report TR-12-010, September 2012.

B. Raghavan, T. Koponen, M. Casado, A. Ghodsi, S. Ratnasamy, and S. Shenker. Software-Defined Internet Architecture. Proceedings of the ACM Workshop on Hot Topics in Networks (HotNets-XII), Seattle, Washington, October 2012.

B. Raghavan, T. Koponen, A. Ghodsi, V. Brajkovic, and S. Shenker. Making the Internet More Evolvable. ICSI Technical Report TR-12-011, September 2012.

H. O. Song, S. Zickler, T. Althoff, R. Girshick, C. Geyer, M. Fritz, P. Felzenszwalb, and T. Darrell. Sparselet Models for Efficient Multiclass Object Detection. Proceedings of the 12th European Conference on Computer Vision (ECCV 2012), pp. 802-815, Firenze, Italy, October 2012.

R. M. Stern and N. Morgan. Hearing is Believing: Biologically-Inspired Feature Extraction for Robust Automatic Speech Recognition. Signal Processing Magazine, Vol. 29, No. 6, pp. 34-43, November 2012.

N. Tandon, D. Rajagopal, and G. de Melo. Markov Chains for Robust Graph-Based Commonsense Information Extraction. Demo at the 24th International Conference on Computational Linguistics (COLING), Mumbai, India, December 2012.

A. Topfer, O. Zagordi, S. Prabhakaran, V. Roth, E. Halperin, and N. Beerenwinkel. Probabilistic Inference of Viral Quasispecies Subject to Recombination. Journal of Computational Biology, Vol. 20, No. 2, pp. 113-123, February 2013.

TJ Tsai and N. Morgan. Longer Features: They Do a Speech Detector Good. Proceedings of the 13th Annual Conference of the International Speech Communication Association (InterSpeech 2012), Portland, Oregon, September 2012.

O. Vinyals, Y. Jia, L. Deng, and T. Darrell. Learning with Recursive Perceptual Representations. Presented at the 26th Annual Conference on Neural Information Processing Systems (NIPS 2012), Lake Tahoe, Nevada, December 2012.

A. Wundsam and M.Yu. NOSIX: A Portable Switch Interface for the Network Operating System. ICSI Technical Report TR-12-013, October 2012.

W. Yang, J. Novembre, E. Eskin, and E. Halperin. A Model-Based Approach for Analysis of Spatial Structure in Genetic Data. Presented at the Annual Meeting of the American Society of Human Genetics, San Francisco, California, November 2012.

S. X. Yu and D. Lisin. Change Blindness Reveals Distinctive Featural Processing in Size, Color, and Orientation. ICSI Technical Report TR-13-001, February 2013.

multimedia analysis, continued

continued from page 3

students understand how photos, videos, and text updates shared on social media can be used against them. Much of this work is in collaboration with Dan Garcia from UC Berkeley and the Berkeley Foundation for Opportunities in Information Technology, an ICSI project that supports female students and students from underrepresented ethnicities who want to enter a career in computer science.

In that project and others, researchers are interested in chains of inference, or the aggregation of public and seemingly innocuous information from different Web sites in order to attack privacy. For instance, the researchers have shown how it is possible to link accounts on different sites – such as Yelp, Twitter, and Flickr – based on the length of videos posted and other factors, even when the username is different.

It was the privacy implications of such innocuous public information that first sparked Friedland's interest in multimedia analysis. In 2008, Friedland and Sommer published a technical report that coined the term "cybercasing," or the use of information freely available online to mount real-world attacks. The case studies they presented relied on geo-tags, detailed information about where a video or photo was taken that is automatically embedded in media captured by devices such as smartphones. They showed how easy it is, for example, to extract the geo-tags of a photo in a Craigslist ad selling valuables and to use information about the best time to call to infer when the seller was away from home.

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.

INTERNATIONAL
COMPUTER
SCIENCE
INSTITUTE

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

facebook.com/ICSIatBerkeley

twitter.com/ICSIatBerkeley