

INTERNATIONAL
COMPUTER SCIENCE
INSTITUTE

10100101
101CS101
10100101

ANNUAL REPORT
TWO THOUSAND
AND TWELVE

LETTER FROM DIRECTOR ROBERTO PIERACCINI

SMALL LAB, BIG IDEAS, DEEP UNDERSTANDING

ICSI is one of the few independent institutes in the world where the culture of curiosity-driven, unfettered research is preserved, encouraged, and cherished. Because of that, several world-class scientists choose to make ICSI their home or to maintain close scientific ties with it. During my first year at ICSI as its CEO and director, I saw the Institute grow stronger in several areas that were already well established. Internet networking and security, speech, multimedia, and artificial intelligence research achieved important research goals, and were validated by equally important new projects sponsored by various US government agencies, such as NSF and IARPA. Vision research kept moving ahead. Scientifically significant efforts in algorithms, computational biology, brain networks, and computer architecture research achieved important milestones. ICSI continues to be at the forefront of the scientific community in all its fields by pursuing research on big ideas, and by trying to achieve deep understanding of the most crucial questions of computer science. What's the new Internet architecture like? What are the forces that shape the Internet underground economy and crime? Why do we seem to be unable to achieve human performance with speech recognition as soon as the acoustic conditions deviate from optimal, no matter what we do? Do we really need a lot of data to recognize a language? Can a computer visually discriminate, with the accuracy of a human expert, among subcategories of objects, like different models of cars of the same make, or different species of birds

of the same genus? Can we design efficient algorithms to identify concepts in millions of videos, like feeding an animal or changing a tire, without relying on textual tags? Can we automatically identify and comprehend metaphors and come to a deeper understanding of language? Where does spam come from, and what are its socio-economic foundations? Deeply understanding the implications of these big questions, and attempting to find answers to them, is what researchers at ICSI are passionate about. But there is more to that than just individual curiosity.

Turing Award laureate Richard Hamming, a collaborator of Claude Shannon at Bell Labs and the founder of ACM, talked about great research. He said that a great researcher, not just a good one, is one who chooses to work with passion and dedication on important problems. Although there are other qualities that make a great researcher, as Hamming points out, the choice of which problem to attack is a crucial one. And I believe that ICSI, a small lab compared with other similar institutions, has great researchers that have the courage and the determination to attack big and important problems.





TABLE OF CONTENTS

| | |
|------------------------|----|
| ABOUT ICSI | 2 |
| BOARD OF TRUSTEES | 6 |
| NETWORKING GROUP | 8 |
| SPEECH GROUP | 12 |
| AI GROUP | 18 |
| VISION GROUP | 22 |
| ALGORITHMS GROUP | 26 |
| ARCHITECTURE GROUP | 28 |
| INTERNATIONAL PROGRAMS | 30 |
| TECHNICAL CONTENT | 38 |
| NEWS | 70 |

The International Computer Science Institute (ICSI) is a leading center for research in computer science. As one of the few independent, nonprofit research institutes in the United States, the Institute provides a venue in which academic, industrial, and government interests inspire world-class research and ground-breaking developments.

Although the Institute defines six major areas of research—in 2012, they were artificial intelligence, computational algorithms, computer architecture, networking, speech, and vision—research directions are largely determined by the interests of the principal investigators, who span the spectrum of academic disciplines from anthropology to mathematics. Researchers bring the latest developments in psychology to bear on computer interfaces, use algorithms to map our biological processes, and work toward multi-modal computer understanding in an environment that encourages close collaboration between divergent academic areas. The common threads are originality, depth of experience, and excellence.

The international program has been an integral part of ICSI since its inception. ICSI was founded in 1986 and officially inaugurated in 1988 as a joint project of the Electrical Engineering and Computer Science Department (and particularly of the Computer Science Division) of UC Berkeley and the GMD, the Research Center for Information Technology GmbH in Germany. Today, visitors travel to ICSI, both through its active sponsored visitor agreements with Finland and Germany and in order to work on specific projects.

There are approximately eighty scientists in residence at ICSI including principal investigators, postdoctoral Fellows, visiting researchers, and students. The CEO of the Institute is Dr. Roberto Pieraccini.

AFFILIATION WITH UC BERKELEY


Since its founding, the Institute has maintained an affiliation with the University of California at Berkeley. Several of its scientists hold joint faculty appointments at UC Berkeley, teaching graduate and undergraduate courses and supervising students who pursue their doctoral thesis research at the Institute. Its offices are located in downtown Berkeley, adjacent to the UC Berkeley campus and in close proximity to San Francisco and the Silicon Valley.

ICSI FUNDING

ICSI is sponsored by a range of US federal, international, industrial sources, and two family foundations. The figure below gives the relative distribution of funding among these different sponsoring mechanisms.

US federal funding in 2012 came from a range of grants to support research institute-wide. Most of this funding came from the National Science Foundation, IARPA, and DARPA, with other Federal funds coming from AFRL, ARL, California Connects, DOC, DOE, NIH, and ONR.

International support of our visitor program came from the Ministry of Education and Research (via the DAAD) in Germany, TEKES (via the Aalto University in Finland), and Nanyang Technical University (NTU) in Singapore. Industrial support was provided by BT-CITRIS, Cisco, Comcast, Google, Intel, Microsoft, Qualcomm, TASC and Toyota. Foundation support was provided by the S. D. Bechtel, Jr. Foundation and the John Templeton Foundation.



Total ICSI revenue was \$12.5M in 2012.



ABOUT ICSI

ICSI MANAGEMENT



Roberto Pieraccini
Chief Executive Officer



Nelson Morgan
Deputy Director



Scott Shenker
Chief Scientist



Theresa Hogg
Chief Financial Officer



David Johnson
Systems Manager



Maria Eugenia Quintana
Operations & Grants Manager

ICSI SPONSORS

Government sponsors



International sponsors



Private sponsors



Industry sponsors



ABOUT ICSI

ICSI RESEARCH PARTNERS & COLLABORATORS

Raytheon
BBN Technologies

**Carnegie
Mellon
University**

 **CASE WESTERN RESERVE
UNIVERSITY** EST. 1826

 **COLUMBIA UNIVERSITY**
IN THE CITY OF NEW YORK

 **cpacket™**
NETWORKS

 **cyCorp**

 **DECISIVE
ANALYTICS**
Solutions For Success

**Georgia
Tech**

 **HELSINKI
INSTITUTE FOR
INFORMATION
TECHNOLOGY**

IBM®

 **idiap**
RESEARCH INSTITUTE

 **Kitware**

 **BERKELEY LAB**

 **Micron**

 **Massachusetts
Institute of
Technology**

 **NCSA**

 **NORTHWESTERN
UNIVERSITY**

**THE
OHIO
STATE
UNIVERSITY**

 **PRINCETON
UNIVERSITY**

Reservoir Labs

 **SARNOFF®
Corporation**

 **専修大学
SENSHU UNIVERSITY**

SIEMENS
corporate research

SRI International

**Stanford
University**

 **TEL AVIV UNIVERSITY**

 **UNIVERSITY OF CENTRAL FLORIDA**

 **THE UNIVERSITY OF
CHICAGO**

 **THE UNIVERSITY
OF EDINBURGH**

 **UNIVERSITY OF HELSINKI**

 **UMASS**

 **UNIVERSITY OF
MICHIGAN**

 **USF UNIVERSITY OF
SOUTH FLORIDA**

**THE UNIVERSITY OF
TEXAS
AT AUSTIN**

 **UT DALLAS**

W

 **Penn**
UNIVERSITY OF PENNSYLVANIA

Berkeley
UNIVERSITY OF CALIFORNIA

UCIrvine
University of California, Irvine

UCMERCED
UNIVERSITY OF CALIFORNIA

UC San Diego

UCLA

UCSB

UCSF

 **USC University of
Southern California**

VASSAR

BOARD OF TRUSTEES



BOARD OF TRUSTEES



Greg Badros, 2012-
Facebook




Hervé Bourliard, 2001-
IDIAP and EPFL, Switzerland



Deborah Crawford, 2012-
Drexel University



David Culler, ex officio
UC Berkeley



Graham Fleming, ex officio
UC Berkeley



Mazin Gilbert, 2012-
AT&T Labs



Barbara Grosz, 2012-
Harvard University



Clifford Higgerson, 2001-
Walden International



Eric Horvitz, 2012-
Microsoft Research



Martti Mäntylä, 2012-
Aalto University



Nelson Morgan, 1999-
ICSI



David Nagel, 2003-
Ascona Group



Peter Norvig, 2009-2012
Google



Roberto Pieraccini, ex officio
ICSI




Prabhakar Raghavan, 2007-
Google



Scott Shenker, 2003-
UC Berkeley and ICSI




David Tennenhouse, 2007-2012
New Venture Partners



Wolfgang Wahlster, 2001-
DFKI, Germany

NETWORKING





NETWORKING

The last decade has seen dramatic transformations in both the capability and use of malicious software. Among the scourges now routinely inflicted by Internet attackers are distributed denial-of-service attacks, widespread information theft, click fraud, spamming, and phishing. ICSI pursues research spanning many dimensions of this important problem space, ranging from technologies to detect attacks and analyze their aftermath to approaches for measuring, understanding, and disrupting attacker operations and goals.

PROJECTS

001 – BRO

A powerful, high-performance platform for analyzing network activity in rich semantic depth. Bro provides the user with a domain-specific scripting language used to define real-time analysis of network traffic in terms of events reflecting different elements of the traffic's activity.

010 – NETALYZR

A network troubleshooting service that allows anyone with a Java-enabled Web browser to evaluate their Internet connectivity's current status, contributing in turn to a wide-ranging measurement effort studying the Internet's health. To date, it has been run more than 800,000 times by more than half a million Internet systems, providing a wealth of data from an otherwise rarely explored perspective.

011 – SOFTWARE-DEFINED NETWORKING

A new approach to network control planes that was invented at ICSI jointly with colleagues at Stanford and later Nicira. While a relatively recent development, with the first full-blown realization in 2008 comprising the NOX network operating system and the OpenFlow forwarding model, SDN has already gathered wide support from industry.



100 – FUTURE INTERNET ARCHITECTURE

Work on the broad issue of fostering architectural innovation to the Internet, with the goal of allowing the Internet to evolve over time and adapt to future requirements. The key design challenge is to identify the absolutely minimal portion of the design that needs to be universally agreed upon (and therefore fixed for long periods of time) and use this core as a foundational framework that allows the rest of the architecture to evolve more freely.

101 – CESR

An interdisciplinary effort that takes the view that, while security is a phenomenon mediated by the technical workings of computers and networks, it is ultimately a conflict driven by economic and social issues that merit a commensurate level of scrutiny. Begun in 2012 and initially supported by a five-year grant from the National Science Foundation, CESR has a primary goal of tackling key social and economic elements of security: how the motivations and interactions of attackers, defenders, and users shape the threats we face, how they evolve over time, and how they can best be addressed.

PEOPLE

RESEARCHERS

Scott Shenker
Mark Allman
Vjekoslav Brajkovic
Chris Grier
Seth Hall
Christian Kreibich
Barath Raghavan
Robin Sommer
Nick Weaver

VISITORS

Bernhard Amann (Germany)
Haixin Duan
Oana Goga
Dorgival Guedes (Brazil)
Eemil Lagerspetz (Finland)
Christof Leng (Germany)
Sasu Tarkoma (Finland)
Renata Teixeira
Andreas Wundsam (Germany)
Daniel Warneke (Germany)
Jiao Zhang

STUDENTS AND INTERNS

Soumya Basu
Jason Croft
Lorenzo DeCarli
Mohan Dhawan
Igor Ganichev
Dina Hadziosmanovic
Shuang Hao
Jian Jiang
David Kantola
Sheharbano Khattak

Junaid Khalid
Nupur Kothari
Paul Pearce
Justin Samuel
Colin Scott
Kristin Stephens
Amin Tootoonchian
Matthias Vallentin
Sam Whitlock
Kyriakos Zarifis

NSF AWARDS \$10 MILLION GRANT TO ICSI AND COLLABORATORS TO STUDY HUMAN ELEMENT OF CYBERCRIME

<https://www.icsi.berkeley.edu/icsi/news/2012/09/frontier-cybercrime-economy>

ICSI, the University of California, San Diego, and George Mason University have received a \$10 million, five-year grant from the National Science Foundation to investigate the roles played by economics and social interactions in Internet security.

"Online crime is relentlessly populating a subterranean layer of digital space, during a time when the Internet is playing an increasingly integral and significant part in today's world economy, as well as many facets of our daily lives. Advanced computer science research has a moral obligation and the right tools to investigate this phenomenon and help prevent its dangers," said Director Roberto Pieraccini.

"During our earlier work on analyzing the factors that go into making spam a profitable form of cybercrime, we were deeply struck by the significance of the human side of the equation," said Vern Paxson of ICSI's Networking Group, one of the leaders of the project and a UC Berkeley professor. "Non-technical considerations span business concerns, issues of trust-amongst-thieves, and the rise of social media as both a new domain that cybercrime is expanding into, and a way to track interactions amongst the criminals themselves."

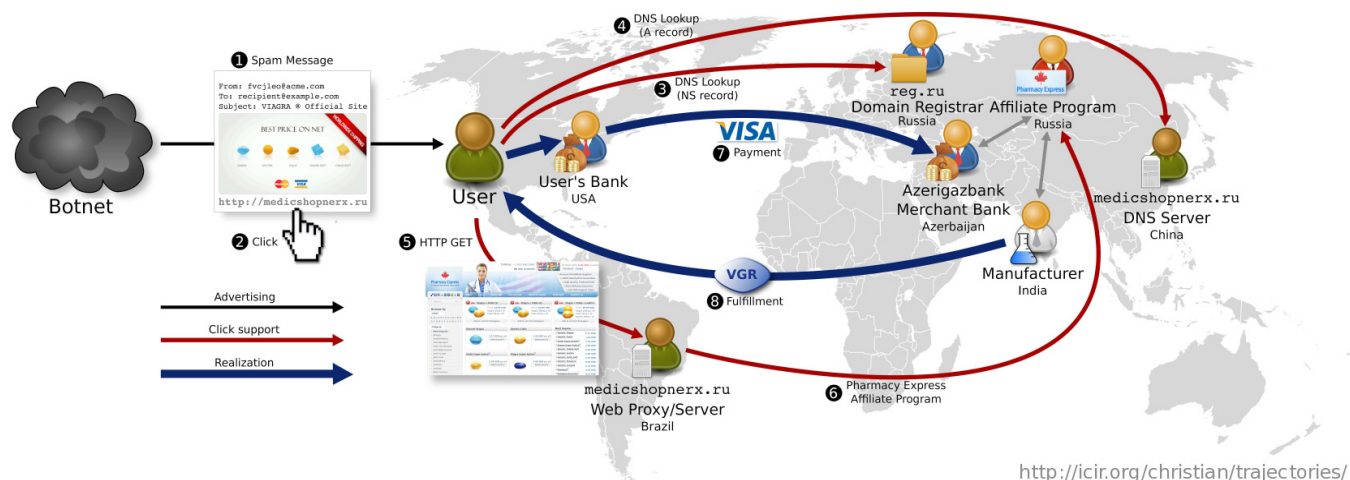
Security research has tended to focus on the technologies that enable and defend against attacks. The new project, led by Paxson and Stefan Savage of UC San Diego, emphasizes the profits that motivate the majority of Internet attacks, the

elaborate marketplaces that support them, and the relationships among cybercriminals, who rely upon each other for services and expertise.

Paxson and Savage have a long history of working together on Internet security. In 2004, they established the NSF-funded Center for Internet Epidemiology and Defenses (CCIED), which they have led since then, and more recently they've led a large-scale effort funded by the Office of Naval Research on infiltrating the "botnets" that attackers often use for their attacks. Recent results from CCIED include findings that just three banks authorized 95 percent of credit card sales of goods advertised through spam. This kind of research aims to help the fight against spam and malware by exposing weak points in the spam economic chain.

The new project, a natural outgrowth of this, focuses on the human element of cybercrime, including how social media such as Facebook and Twitter provide new opportunities for attacks and manipulation. By better understanding the roles that economics and social interactions play in cybercrime, the researchers say, defenders can identify the most effective opportunities for interventions and defenses.

The work will be carried out by the Center for Evidence-Based Security Research, which was established by the NSF grant.





NETWORKING

AN INSIDE VIEW OF THE ONLINE PHARMACEUTICALS INDUSTRY

<https://www.icsi.berkeley.edu/icsi/blog/online-pharmaceutical-industry>

Financial records of three vendors that sell unauthorized and counterfeit pharmaceuticals over the Internet show, among other things, that they rely on a relatively small number of affiliate advertisers to drive traffic to their sites. An analysis of the records by Networking researchers and their collaborators gives a rare insider's view of the finances of illicit online activity.

The study analyzes financial records of GlavMed, SpamIt, and RX-Promotion, which sell erectile dysfunction, male enhancement, and other drugs and which pay third-party affiliates to advertise their Web sites through spam and abusive search engine optimization methods. The records were leaked to journalists, forums, and law-enforcement agencies by the rivaling organizations as part of what Brian Krebs, a network security blogger and a co-author on the paper, calls the Pharma Wars. The records used in the study cover all sales made by GlavMed and SpamIt from 2007 to early 2010, and all sales made by RX-Promotion to U.S. customers during 2010. They covered almost 1.5 million orders totaling \$185 million.

The resulting paper is "perhaps the most detailed analysis yet of the business case for the malicious software and spam epidemics that persist to this day," Krebs wrote on his blog. It's unusual in that it relies on ground truth data – the actual daily finances of an industry that until now has been shrouded in secrecy. The dataset allowed the researchers to corroborate financial estimates they had made using a variety of empirical measurements published in earlier papers.

The paper was presented in August at the USENIX Security Symposium. In addition to Krebs and Networking researchers Christian Kreibich and Nicholas Weaver, its authors include researchers from George Mason University and UC San Diego.

Spam and other illicit advertising form part of what the researchers describe as an underground economy, in which profit is the main motivating factor behind Internet attacks and scourges such as email spam. In recent years, the underground economy has experienced a division of labor, with different tasks performed by different organizations, which get a cut

of the sales. In this "affiliate program" model, a sponsor pays third-party affiliates on commission to drive traffic to its Web site. Spam is one major method of doing this. Sponsors also partner with third parties to process payments.

The study found that 10 percent of affiliate advertisers account for more than three-quarters of revenue, and that just three payment service providers are responsible for 84 percent of revenue. The researchers say these may be weak points in the spam economic pipeline, vulnerable to counterattacks.

The analysis revealed some other interesting facts. Even though the number of orders fulfilled by GlavMed and SpamIt dropped beginning in mid-2009, the number of new customers remained steady throughout the period covered by the data, suggesting that the market for online pharmaceuticals is far from saturated.

Repeat customers constituted more than a quarter of GlavMed's and SpamIt's revenue. The most popular drugs sold by the two were related to erectile dysfunction, which accounted for 75 and 82 percent, respectively, of their revenue.

RX-Promotion, on the other hand, relied more heavily on sales of controlled drugs that are potentially habit-forming, such as pain medications and opiates. While only 14 percent of orders fell in this category, they represented 32 percent of total revenue. These drugs are generally more closely regulated due to their addictive nature.

The work also found that because of high commissions to advertisers, who receive between 30 and 40 percent of each sale that they facilitate, and other costs, such as those for shipping and processing credit cards, the program operators receive less than 20 percent of sales revenue.

Related Paper: "PharmaLeaks: Understanding the Business of Online Pharmaceutical Affiliate Programs." D. McCoy, A. Pitsillidis, G. Jordan, N. Weaver, C. Kreibich, B. Krebs, G. M. Voelker, S. Savage, and K. Levchenko. Proceedings of the 21st USENIX Security Symposium, Bellevue, Washington, August 2012. Available online at http://www.icsi.berkeley.edu/icsi/publication_details?n=3314

SPEECH



SPEECH

The Speech Group conducts research in the areas of algorithms, architectures, and systems for speech and audio signal processing and pattern recognition. In the 1990s the emphasis of this group's work gradually became the application of these areas of study to problems in spoken language processing, and particularly in robust automatic speech recognition. Key research areas of the Speech Group include auditory-inspired signal processing, statistical modeling, rich transcription of natural and impromptu meetings, and rich transcription of broadcast and conversational telephone speech.

PROJECTS

001 - OUCH

Quantitative characterization of speech recognition errors. A large number of the errors in modern speech recognizers are due to the inaccuracy of the conditional independence assumptions in HMMs. Other sources of errors are the mismatch between training and test conditions, and the signal representation. This study produces tools that are useful for future advances in ASR by providing more detailed diagnostics beyond simple word error rate.



010 - SWORDFISH

Speech recognition in under-represented languages. This project develops tools for the generation of "language packs" for any written language, without large amounts of data.

011 - MULTIMODAL LOCATION ESTIMATION

A collaboration with UC Berkeley faculty member Kannan Ramchandran. Location estimation is the task of estimating where digital media was recorded. This project aims to use the GPS-tagged media available on the Web as a training set for an automatic location estimator.

PEOPLE

RESEARCHERS

Nelson Morgan
Jaeyoung Choi
Dan Ellis
Gerald Friedland
Luke Gottlieb
Jim Hieronymus
Adam Janin
Howard Lei
Benjamin Martinez Elizalde
Nikki Mirghafori
Hari Parthasarathi
Roberto Pieraccini
Steve Wegmann
Chuck Wooters

VISITORS

Hai Do (Singapore)
Seppo Enarvi (Finland)
Arlo Faria (Brazil)
Daniel Göhring (Germany)
Frantisek Grezl
David Imseng (Switzerland)
Mikko Kurimo (Finland)
Kalle Palomaeki (Finland)
Korbinian Riedhammer (Germany)
David Suendermann

STUDENTS AND INTERNS

Sarah Berson
Shuo-Yiin Chang
James Fung
Mary Knox
Martin Ortega
Yun Park
Suman Ravuri
Lara Stoll
TJ Tsai
Oriol Vinyals

ICSI LEADS TEAM RESEARCHING WAYS TO BUILD SPEECH RECOGNITION SYSTEMS FOR NEW LANGUAGES UNDER SEVERE DATA AND TIME CONSTRAINTS

<https://www.icsi.berkeley.edu/icsi/news/2012/11/swordfish>

ICSI is leading a research team under the IARPA Babel Program that is focused on 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 for development of speech models for languages that could emerge as important in the future.

"The goal of the Babel program is to rapidly build speech recognition systems to support effective keyword search for new languages using limited amounts of transcribed speech recorded in real-world conditions," said Mary Harper, the IARPA Program Manager in charge of the Babel program.

Using only a fraction of the training data usually required, the team aims to build speech recognition systems for several languages in just one week by the end of the program.

"ICSI excels at intellectual challenges and unique approaches to research. This is an intriguing project that puts significant constraints on our researchers as a means to discover better ways to develop automatic speech recognition systems," said Roberto Pieraccini, director and president of ICSI.

By working on a variety of languages with time and data restrictions, the team will research basic principles of speech technology rather than incremental improvements to existing technology. In addition, this research will be useful in enabling keyword-search systems for those languages that do not have large amounts of transcribed audio.

"The speech recognition systems we've built in the past have the curse of being reasonably good, particularly for a few

languages and speech recorded in good acoustic conditions, which has often reduced the impetus to significantly change the technology," said Professor Nelson Morgan, deputy director and leader of the Speech Group at ICSI. "This project strongly pushes us to solve fundamental problems in speech recognition to address the Babel challenge."

In each of the four periods of the project, named the SWORDFISH project, the team will be given a set of languages and will be tasked with developing methods to quickly build a system. Speech recognition systems are typically trained on thousands of hours of transcribed audio. In this project, the team was initially given only 80 hours of conversational speech for each language, and in each succeeding period a smaller fraction of the audio is transcribed. At the end of each period, the team will be given a new language to build a system – initially in four weeks, but by the end of the program down to just one week.

In addition to Morgan, the leaders of the team are Steven Wegmann of ICSI, Professor Mari Ostendorf of the University of Washington, Professor Janet Pierrehumbert of Northwestern University, Professor Eric Fosler-Lussier of The Ohio State University, and Professor Dan Ellis of Columbia University. Morgan says an important element of the project is that these team leaders have had strong previous research ties with one another in research topics that are essential to the Babel problem.

The project is funded by the Intelligence Advanced Research Projects Activity (IARPA), a research arm of the Office of the Director of National Intelligence, which invests in high-risk/high-payoff research programs.

ESTIMATING WHERE FLICKR VIDEOS WERE SHOT

<https://www.icsi.berkeley.edu/icsi/blog/estimating-video-location>

ICSI researchers and their collaborators have developed an algorithm that estimates where random Internet videos were shot by training itself on both videos for which locations have already been identified and those for which they haven't. It's the first multimodal location estimation method to use both training and test data to improve its results and, in some cases, is over 10 percent more accurate than existing methods. The improvement is particularly high when few training videos are available.

The work was performed on videos selected at random from Flickr but could be applied to videos uploaded to other social media sites such as YouTube and Twitter. A system capable of accurately placing such media would allow users to quickly and easily group related videos in large collections, for example, and might also improve performance-based services.

In location estimation of videos, data from videos whose locations are known is used to develop profiles of their respective locations. This data may include text data such as location tags, visual cues such as textures and colors, and sounds such as bird song. The attributes of a test video are then compared against these profiles and its location is estimated.


However, a new method developed by ICSI researchers and collaborators also compares test videos to each other. Say, for example, that a training set included a video with the tags {berkeley, california}. Existing methods would easily be able to determine where a video with the tags {berkeley, sathergate, california} was shot, but unable to determine where a video with the tags {sathergate, california} was shot since it was not tagged {berkeley}. The new algorithm looks at similarities between the two test videos – in this case, the tags sathergate and california – and applies what it knows about the first to the second.

After the system estimates the locations of the videos, it groups videos by region and analyzes their auditory and visual features to further refine its estimate.

This approach allows researchers to more accurately determine locations when little data is available since test videos are essentially treated as training data. Training data can only comprise photos and videos with known locations – and that's only about 5 percent of the media uploaded to the Internet, with even less data available for certain regions such as South Africa. By taking advantage of the huge amount of data whose location is unknown, the approach allowed for more than 10 percent improvement in some cases over existing methods.

The results are part of ICSI's ongoing work in multimodal location estimation, the placing of photos and videos using a variety of indicators such as the text in tags, sounds in audio tracks, and visual cues. In one study, for example, ICSI researchers analyzed

the sounds of ambulances from different cities around the world, allowing them to place videos that included an ambulance. Similar work has been done using visual cues such as texture and color.



Related Paper: "Multimodal Location Estimation of Consumer

Media: Dealing with Sparse Training Data." Jaeyoung Choi, Gerald Friedland, Venkatesan Ekambaram, and Kannan Ramchandran. Proceedings of the IEEE International Conference on Multimedia and Expo, Melbourne, Australia, July 2012. Available online at <http://www.icsi.berkeley.edu/cgi-bin/pubs/publication.pl?ID=3274>.

QUALIFYING MECHANICAL TURK USERS FOR SKILLED TASKS

<https://www.icsi.berkeley.edu/icsi/blog/mechanical-turk>

Mechanical Turk, Amazon's popular crowdsourcing platform, is used by many people and organizations who need repetitive tasks completed. For rapid completion of unskilled (easy) tasks online, crowdsourcing can be an easy and affordable solution.

But what if you need people to complete a skilled task? How do you filter the crowd to find qualified people? Researchers at ICSI decided to find out, and lead author Luke Gottlieb presented their results at the ACM Multimedia conference which was held in October in Nara, Japan.



In their experiment, they tried to find people who were qualified to determine the location of video clips they were sent. ICSI is active in multimedia and geolocation research, so the ability to compare the accuracy of skilled humans to that of machine geolocation for video will be useful. The experiment, in addition to providing valuable insight into the potential use of Mechanical Turk for skilled tasks, provides a benchmark for how well skilled humans can perform geolocation compared with the ICSI Multimodal Location Estimation (MMLE) system.

Prior to screening users on Mechanical Turk, the team presented video clips to other researchers to narrow down a test set of clips they believed could be accurately located within five minutes by a skilled human.

In the qualifying test, it was important that the videos didn't require specific cultural knowledge, and that it would be an entertaining enough task for many users to want to try it. The researchers set a standard of 80 percent accuracy in order for a user to qualify as skilled for the task, and provided detailed instructions for completing the task. They found that 20 percent of Mechanical Turk users who tried the test set proved to be skilled enough for the task.

Once the qualified users were selected from the test set, the pay rate was increased so that users would still be motivated to complete the tasks even if they weren't quite as entertaining as the tasks in the training set. Keeping users motivated is an important part of crowdsourcing, in order to have enough of a crowd of users who want to complete the tasks. By keeping the tasks entertaining in some way, enough users will complete the tasks despite a relatively low pay rate. This enables researchers to recruit a large team of skilled users affordably.

A paper on the process was presented at ACM Multimedia, and Gottlieb and his colleagues have been invited to write an expanded follow-up article for the journal IEEE Transactions on Multimedia.

MMLE Web page:

<http://mmle.icsi.berkeley.edu/mmle/index.html>

Related Paper: "Pushing the Limits of Mechanical Turk: Qualifying the Crowd for Video Geo-Location." Luke Gottlieb, Jaeyoung Choi, Pascal Kelm, Thomas Sikora, and Gerald Friedland. 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.

RESEARCH PROJECT SEEKS TO IDENTIFY UNDERLYING CHALLENGES TO CURRENT ASR TECHNOLOGY

<https://www.icsi.berkeley.edu/icsi/news/2012/06/ouch-press-release>

A new research project at ICSI is focused on exploring automatic speech recognition (ASR) to understand the limitations and challenges from current technologies. Sponsored by the Intelligence Advanced Research Projects Activity (IARPA) via the Air Force Research Lab (AFRL), the research aims to use its conclusions to lead to new methods for improving ASR technology.

"This is a unique research project in that we are qualitatively and quantitatively exploring what is wrong with automatic speech recognition. From that we hope to gain insights into how we can improve ASR, potentially going forward in entirely new directions," said ICSI Deputy Director Nelson Morgan, who leads the project. "When you don't know specifically what is wrong with a technology, you are left with a hit-or-miss situation. This research should give us some clarity." Morgan is also the leader of the Speech Group.


The research project, named Project OUCH, includes two major parts. The first is an in-depth look at the assumptions behind acoustic modeling, a key component of ASR that creates statistical representations of each of the distinctive sounds that make up words. This will enable ICSI researchers to discover technical challenges that prevent ASR from being more accurate.

The second part is a broad survey of experts and colleagues in the field, asking for perceptions on where ASR technology is effective, where it fails, and what its shortcomings are. This study will include interviews of practitioners and a review of recent literature to derive community consensus on what approaches don't work and to develop guidelines for future analysis.

Steven Wegmann and Jordan Cohen serve as co-principal investigators of the research. Wegmann oversees the in-depth acoustic modeling phase, and Cohen the breadth field survey phase.

This research is supported by the Intelligence Advanced Research Projects Activity (IARPA) via Air Force Research Laboratory (AFRL) contract number FA8650-12-C-7217. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright annotation thereon. Disclaimer: The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of IARPA, AFRL, or the U.S. Government.

A1



Artificial intelligence research is carried out by a highly interdisciplinary group of computer and cognitive scientists, linguists, psychologists, and neuroscientists investigating the cognitive and neural underpinnings of language and thought. Stemming from this basic research thrust, the group is also active in applying the insights gained in a variety of applications, from natural language understanding, machine learning and inference, semantic resource development, and education and development in underserved communities.


PROJECTS

001 – METANET

A new major multi-university effort sponsored by IARPA to investigate the relationship between metaphor, computation, culture, and cognition.

010 – DIGITAL FUNCTIONING IN UNDERSERVED POPULATIONS

A new project combining theories from cognitive science on framing and metaphor and social science findings and techniques to evaluate methods to improve the adoption and use of information technologies by underserved populations. Funded by the Department of Commerce's National Telecommunications and Information Administration through its Broadband Technology Adoption Program (BTAP).



011 – FRAMENET

A lexical database of English that is both human- and machine-readable. FrameNet is gaining importance in the field of computational linguistics as semantics and meaning are becoming increasingly relevant in NLP applications.

100 – HUMAN- COMPUTER INTERACTION

Building meaningful natural language-based interaction systems by applying theories and computational models of cognition developed at ICSI. Sponsored by a basic research grant from the Office of Naval Research and a Templeton Foundation grant.

PEOPLE

RESEARCHERS

Srini Narayanan
Collin Baker
Julia Bernd
Richard Cook
Jerry Feldman
Charles Fillmore
Luca Gilardi
Noelia González-Cámara
Blanca Gordo
Paul Kay
Dan Klein
George Lakoff
Behrang Mohit
Maria Carmen Penaranda
Miriam Petruck
Gerardo Sanchez

Ekaterina Shutova
Eve Sweetser

VISITORS

Oliver Culo (Germany)
Sergio Guadarrama
Emanuel Kitzelmann (Germany)
Gerard de Melo (Germany)
Hiroaki Sato
Malte Schilling (Germany)

STUDENTS AND INTERNS

Shahrazad Ahmadpour
Mohit Bansal
Jalal Buckley
Ryan Culkin
Elena Gasparini
Joseph Giroux
Elkin Darío Gutiérrez
Jisup Hong
Stacy Hsueh
Sanam Janamian
Patrick Keenan
Jenny Lederer
Russell Lee-Goldman

Christine Leper
Anh Mai
Omoju Miller

Aria Pakatchi
Ayden Parish
Georgy Ramonov
William Sultan
Brandon Thai
Chris Xie
Xihao Zhu

NEW RESEARCH: ALIGNING LEXICAL RESOURCES

<https://www.icsi.berkeley.edu/icsi/blog/aligning-lexical-resources>

ICSI researchers and their collaborators have developed a way to statistically compare two of the most widely used lexical resources for English, FrameNet and WordNet. With the Expected Jaccard Index, researchers can for the first time empirically quantify agreement between WordNet and FrameNet expert annotations of words and phrases pulled from real-world texts. Comparing the annotations helps researchers find problematic gaps in the data and align the resources so they can be used together.

The work is part of an effort to understand how lexical resources differ through annotation of the Manually Annotated SubCorpus (MASC), a subset of the American National Corpus, using FrameNet, WordNet, and other lexical resources. The FrameNet Project, established by Professor Charles Fillmore, housed at ICSI, and led by Collin Baker, defines words through the semantic roles they play and the frames (types of event, relation, or entity) they evoke. WordNet, on the other hand, clusters partially synonymous words into "synsets"; researchers then describe the relationships between the synonym sets. Unlike a frame, which may include different parts of speech (such as verbs and nouns) and words with contradictory definitions (such as antonyms related to the same idea), a single synonym set comprises only synonyms of the same part of speech.

WordNet's information about a word based on its synonyms complements FrameNet's syntactic information about the role it plays in a sentence; the resources, however, do not always align nicely and sometimes define a different set of senses of the same word. The word *curious*, for example, has three senses in FrameNet and only two in WordNet: unlike WordNet, FrameNet distinguishes curiosity as a character trait from curiosity as a mental state.

At the International Conference on Language Resources and Evaluation in May, researchers will present a new statistical measure that shows where WordNet and FrameNet agree well on the meanings of words and phrases, and where they do not. One team at Vassar and Columbia Universities used WordNet to annotate MASC sentences in which a particular word appeared, and another team at ICSI annotated the same sentences with frames and lexical units. Each instance of the word was assigned to a cell in a contingency table, with its WordNet sense on

one axis and the frame it evokes on the other. For example, of the 67 annotated sentences that included the word *curious*, 48 both evoked the typicality frame (FN1) and used *curious* in the WordNet sense of "beyond or deviating from the usual or expected" (WN1). As the researchers expected, *curious* never simultaneously evoked the typicality frame and the WordNet sense "eager to investigate."

If WordNet and FrameNet were perfectly aligned, all of the sentences that evoke one frame would use one WordNet sense, and vice versa. However, that is not always the case. The researchers devised a statistical measure based on the standard Jaccard similarity coefficient to measure how closely aligned the resources are. The measure, the Expected Jaccard Index, will produce a high number when a word's WordNet senses align well with its FrameNet frames – that is, when most of the sentences assigned to one frame have the same WordNet sense in its contingency matrix. The index will produce a low number when sentences with one WordNet sense evoke several frames, or sentences evoking one frame use several WordNet senses.

Some results of the statistical measurement were surprising. Researchers expected to find that the more meanings a word has, the less WordNet and FrameNet would agree on its senses. This was not always true: one of the most closely aligned words, *board*, has six relevant senses in both FrameNet and WordNet, while *trace*, one of the least aligned words, has two WordNet and five FrameNet meanings.

The work has also helped researchers identify frames and senses that need adjustment. While aligning the word *curious*, for which only one lexical unit existed at the time ("unorthodox or unexpected," evoking the typicality frame), FrameNet researchers found that they needed to add two lexical units (for the permanent characteristic of being driven to learn and for the temporary state of being inquisitive). Unexpected results – such as sentences with unpredicted combinations of frames and WordNet senses – help researchers identify aspects of their resources that may be confusing or unhelpful.

The Expected Jaccard Index can also be applied to lexical resources other than FrameNet and WordNet as well as different versions of the same lexical resource.

METANET: A MULTILINGUAL METAPHOR REPOSITORY

<http://www.icsi.berkeley.edu/icsi/blog/metanet>

A large team of researchers from ICSI, UC San Diego, University of Southern California, Stanford, and UC Merced is building a computer system capable of understanding metaphors used in American English, Iranian Persian/Farsi, Russian, and Mexican Spanish.

Metaphors describe abstract ideas in concrete terms and include basic notions, like Time Is Money, and phrases based on them, like spending time. Automatic text analyzers typically have difficulty dealing with metaphors, which rely on a set of complex relationships among words and ideas to be intelligible.

By analyzing metaphors in everyday language, the MetaNet project can shed light on the role metaphor plays in how people from different cultural backgrounds make judgments and decisions. Research on metaphors can help uncover the meanings of words and provide insights into worldviews of particular groups, ultimately helping decision makers interact and communicate more effectively on a global stage.

"Metaphors are at the heart of human communication, and they influence and reflect our worldview," said ICSI CEO Roberto Pieraccini. "We hope to build systems that can extract meaning from text more powerfully and accurately than ever before, and at the same time to better our understanding of how various cultures form different beliefs as a result of prevalent metaphors."

The MetaNet project is led by Professor Srinu Narayanan, the Artificial Intelligence Group leader, and comprises researchers with backgrounds in computer science, linguistics, psychology, and cognitive science. The multi-year project is supported by IARPA, a research arm of the Office of the Director of National Intelligence.

The team is using novel machine learning and inference techniques to detect new metaphors in text and understand them. Over the last two decades, the AI Group has pioneered the use of computational techniques for figurative language interpretation. The system-building team works closely with the cognitive linguistics team, which analyzes the relationships between metaphors and between a metaphor's literal and figurative meanings. This cognitive linguistics work is led by Professor George Lakoff of UC Berkeley, one of the founding fathers of cognitive linguistics. The researchers use methods developed in ICSI's FrameNet Project, led by Dr. Collin Baker and Professor Charles Fillmore, to analyze a repository of metaphors. The researchers hope to use unsupervised machine learning for languages like Farsi, for which there is not much annotated text available. In parallel with this work, cognitive psychologists will study the affect associated with metaphors.

"Our approach is particularly challenging because it requires bridging vastly different fields, from cognitive linguistic analysis to machine learning techniques," Narayanan said. "But this cross-disciplinary approach will give us a better overall understanding of language, metaphors, and their influence on culture and perceptions. Ultimately this will help us better understand how the human brain works."

In the program's first year, IARPA asked the teams participating in the program to look at metaphors related to the concept of Governance. The MetaNet team has now built a fully functional repository with a hierarchical relational network of conceptual metaphors that includes more than 1,000 conceptual mappings and associated frames, metaphors, role bindings, and relations. The conceptual metaphors span basic embodied metaphors that are hypothesized to be cross-linguistic as well as language- and culture-specific mappings related to the first-year target domain of Governance, set by IARPA.

VISION



VISION

The Vision Group focuses on solutions to two core computer vision problems – the perception of human expression for multi-modal interaction, and the recognition of object and scene categories – and is interested in applications ranging from mobile phone interfaces to interactive robotics. Awareness of the environment and of a user's presence and/or expression is critical for smart mobile interfaces, and of course for successful real-world action on mobile agents.

PROJECTS

001 – DOMAIN ADAPTATION

Most successful object classification and detection methods rely on classifiers trained on large labeled datasets. However, for domains where labels are limited, simply borrowing labeled data from existing datasets can hurt performance, a phenomenon known as "dataset bias." We propose a general framework for adapting classifiers from "borrowed" data to the target domain using a combination of available labeled and unlabeled examples.

010 – PERCEPTUAL SECURITY

We have investigated new models for cybersecurity that draw on models for human and machine perception. We have studied shortcomings of and improvements to defenses against user interface security attacks, such as clickjacking, from a perceptual perspective and we have defenses against visual phishing attacks, which rely on fooling victims into thinking an attacker-controlled Web page is content from a legitimate source.

011 – REPRESENTATION LEARNING

Our group has focused on four core visual representation learning problems: efficient learning of "deep" models; sharing sparse part descriptors; fine-scale segmentation models; and time-optimized recognition.

100 – ROBOTIC VISION

To perform useful tasks in everyday human environments, robots must be able to both understand and communicate the sensations they experience during haptic interactions with objects. Toward this goal, we augmented the Willow Garage PR2 robot with a pair of SynTouch BioTac sensors to capture rich tactile signals during the execution of four exploratory procedures on 60 household objects.

101 – FINE-GRAINED RECOGNITION

Pose-normalization seeks to align training exemplars, either piecewise by part or globally for the whole object, effectively factoring out differences in pose and in viewing angle. Our effort factorizes the problem of pose-normalization into (i) the localization of semantic parts and (ii) learning an optimal description.



PEOPLE

RESEARCHERS

Trevor Darrell
Ryan Farrell
Ross Girshick
Kate Saenko
Stella Yu

VISITORS

Stefanie Jegelka
Matthias Kirchner (Germany)
Lorenzo Riano
Erik Rodner (Germany)
Seppo Virtanen (Finland)

STUDENTS AND INTERNS

Tim Althoff
Tobias Baumgartner
Jeff Donahue
Judy Hoffman
Allison Janoch
Sergey Karayev

Jonathan Long
Caroline Modic
Evan Shelhamer
Hyun Oh Song
Tom Yeh
Ning Zhang

AFFILIATED

Mario Christoudias
Ashley Eden

IMPROVING COMPUTER VISION WITH POSELETS

<http://www.icsi.berkeley.edu/icsi/blog/poselets>

Computer vision techniques have trouble recognizing subcategories of objects (for example, a vehicle's model type or a bird's species). A new method developed by ICSI researchers improves automatic recognition of subcategories by first warping small areas of photos to account for differences in pose and angle, and then grouping the areas according to their similarities.

The differences between subcategories are often subtle and localized to small portions of an object. At the upcoming IEEE Conference on Computer Vision and Pattern Recognition, ICSI researchers will present their new method based on poselets.

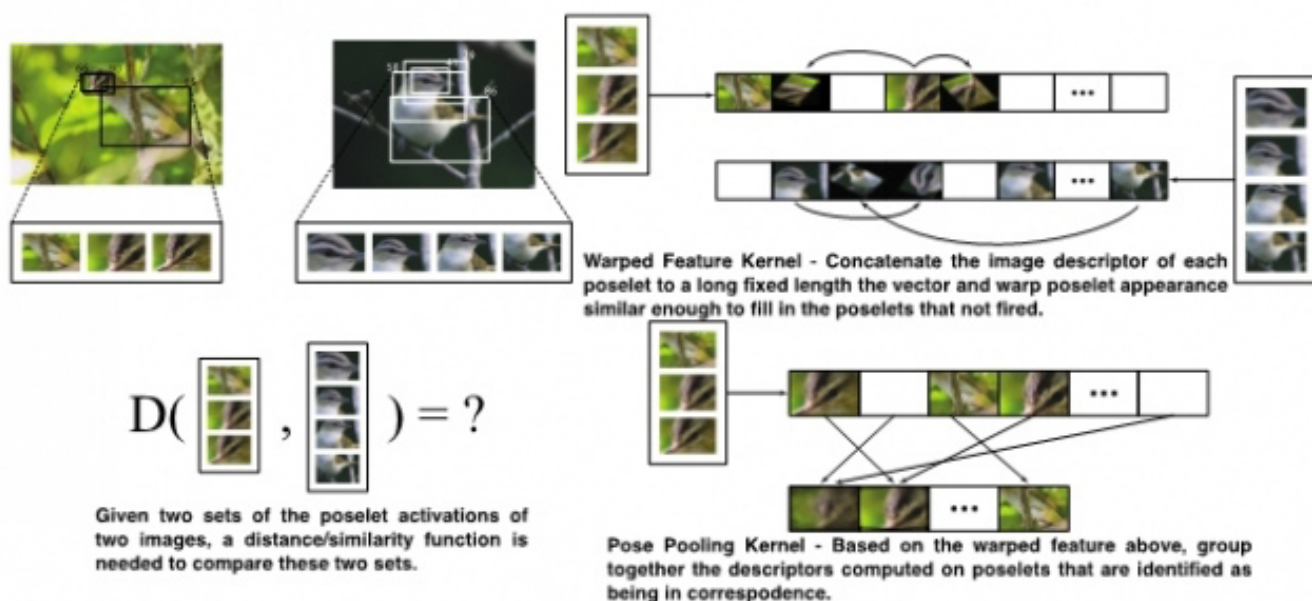
Poselets capture the shape and appearance of a portion of an object, coupling both the object's pose and how it's facing the camera. Each category of object has many poselets, some of which overlap. For example, the category "birds" may have one poselet showing the left side of the bird's face, while another includes both the left side at a different angle as well as the front of the head. For this reason, no single photo has all

possible poselets. The fact that photos generally have different sets of poselets makes comparison between them difficult.

The researchers' new method uses a warped feature kernel to distort one poselet into another where possible. This allows a more direct comparison of two photos that may have different poselets. The new approach also gathers information across poselets – including the distorted poselets – within each image so that they can more easily be compared.

The researchers used these comparisons with nearest-neighbor and kernel-based learning on a set of photos of birds. The method was more effective at identifying subcategories – bird species – than those methods that only consider statistics from the whole image without normalizing poses.

Related Paper: "Pose Pooling Kernels for Sub-Category Recognition." Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, June 2012. Available online at <http://www.icsi.berkeley.edu/cgi-bin/pubs/publication.pl?ID=3300>.



NEW METHOD FOR DETECTING OBJECTS IN IMAGES

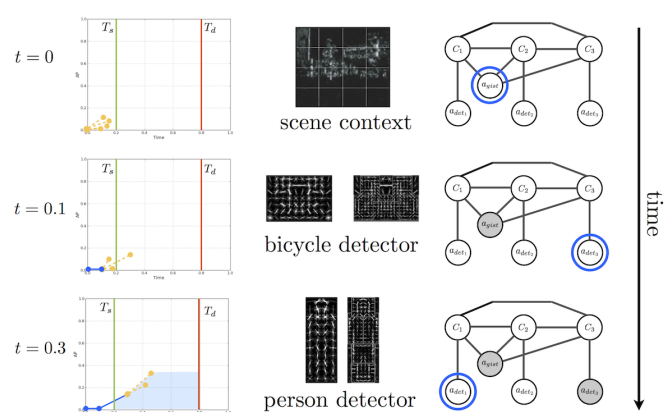
<https://www.icsi.berkeley.edu/icsi/blog/object-recognition>

Using techniques from the field of robotics, Vision researchers and their colleagues have developed a method for detecting objects in images that intelligently selects which object detectors to use and which to ignore in order to complete a task within given time constraints. The paper was presented in December at the Neural Information Processing Systems Conference. It's by Vision researcher Sergey Karayev and group leader Trevor Darrell as well as Tobias Baumgartner of RWTH Aachen University and Mario Fritz of MPI for Informatics, who has worked at ICSI as a postdoc.

The task of visual object recognition is to correctly localize and identify all "objects" in a photograph. In a commonly used computer vision dataset, the PASCAL Visual Object Challenge, objects are labeled by their general category, such as "aeroplane," "car," or "person." For an advertising company on the Internet, the task may be similar – to identify a specific model of car in all images uploaded to a Web site like Flickr or Instagram.

In this and other datasets, certain object classes tend to occur together: for example, buses with cars, or people with bicycles and horses. Current state-of-the-art approaches to detecting all objects of a specific type in an image take about a second to run, per object class. Since many types of objects need to be detected, processing a single image (by running all the object detectors on it) may take well over a minute.

The researchers looked at the case where there isn't enough time to run all the detectors: for example, an advertising company with a long queue of images may only have ten seconds to process an image. In this situation, a subset of object classes to detect needs to be selected so as to maximize the chance of finding the most valuable classes.



This selection of detectors is treated as a sequential decision process. Each detector gives some (imperfect) information about the presence of the corresponding object class in the image, and the method makes use of this information when selecting the next detector.

The decision process is trained using reinforcement (reward-based) learning, a robotics technique that is not often used in computer vision. The reward obtained after selecting a detector is defined as the area under the detection performance versus time curve that the detector contributes to the overall sequence. This results in a system for dynamic selection of detectors whose performance significantly surpasses a static selection baseline.

The bottom line is this: if there are existing object detectors, but not enough time to run all of them in an image, the method can intelligently run the ones that fit within a time budget to maximize the overall multi-class detection performance.

Related Paper: "Timely Object Recognition." Sergey Karayev, Tobias Baumgartner, Mario Fritz, and Trevor Darrell. Proceedings of the 26th Annual Conference on Neural Information Processing Systems (NIPS 2012), Lake Tahoe, Nevada, December 2012.

ALGORITHMS





ALGORITHMS

In recent years, sequencing and genotyping capacity has increased dramatically, leading to exciting discoveries in human genetics regarding the relations between genes and diseases. The extremely large amounts of data generated by these technologies set computational and statistical challenges; a large part of the Algorithms Group's work involves searching for methods that cope with these challenges. The group has been mainly focusing on the development of computational methods for the analysis of high-throughput genomics data, particularly in relation to disease and to population studies. These studies pose many computational problems stemming from the fact that the sequencing and genotyping technologies provide noisy and incomplete data, and since in some cases the information about the genotyped subjects is limited.


PROJECTS

001 - STUDIES OF HUMAN DISEASE

We are not merely developing theoretical methods – we also apply these methods to the analysis of specific diseases. We are collaborating with international groups that study specific diseases, including atherosclerosis, coronary artery disease, leukemia, and non-Hodgkin Lymphoma. In the past year we were able to detect multiple genetic variants that are associated with leukemia and with lipid concentrations, which are correlated with cardiovascular diseases such as atherosclerosis.

010 - METAGENOMICS

The number of bacteria and viruses living within our body is an order of magnitude larger than our own cells. Therefore, in order to study our own genetics and relate the genetic variants to diseases, it is important to characterize the genetic diversity of the microbiome living with us and the variation of bacterial species as a function of disease and other traits. In order to study these bacteria, biologists sample a part of the body and sequence the entire bacterial community. It is then a challenging computational task to separate the sequence data into species. We have developed methods that try to answer this challenge.



011 - POPULATION GENOMICS

Genetic studies of diseases involve the collection and sequencing of many individuals. The analysis of the sequence data involves several computational challenges, including the identification of genetic variations and the assignment of those to specific traits.

100 - BRAIN NETWORKS

We are studying networks derived from brain imaging techniques, such as MRI, in collaboration with researchers at UC San Francisco and the Veterans Administration Medical Center in San Francisco. The goal of these projects is to understand and develop diagnostic procedures for a variety of brain disorders, ranging from traumatic brain injuries, to birth defects and Alzheimer's disease.

101 - MANIPULATION IN DATA CENTERS

This project is developing new algorithms for cloud computing based on economic principles, such as pricing and manipulability. These algorithms are simple to implement, easy to extend, robust to manipulation and provide proper incentives to encourage efficient use of the systems over both short and long timescales.

PEOPLE

RESEARCHERS

Richard Karp
Eric Friedman
Eran Halperin

VISITORS

Michael Elberfeld (Germany)
Itamar Eskin
Farhad Hormozdiari
Ron Shamir
Roded Sharan

STUDENTS AND INTERNS

Robin Gaestel
Luqman Hodgkinson
Inderjit Jutla
Jason Liang
Michael Liang
Graham Tremper

ARCHITECTURE





ARCHITECTURE

The Architecture Group has as its major focus the realization of efficient parallel programmable architectures exploiting advances in circuit and device technologies. The group continues to grow while taking advantage of extensive connections to other research groups in Berkeley and beyond. The primary activities in the Architecture Group at ICSI continue to be in silicon photonics and the vector-thread architecture.

PROJECTS

001 - MONOLITHIC SILICON PHOTONICS

Genetic studies of diseases involve the collection and sequencing of many individuals. The analysis of the sequence data involves several computational challenges, including the identification of genetic variations and the assignment of those to specific traits.



010 - RESILIENT VECTOR-THREAD ARCHITECTURE WORK

We have taped out a further design in 28nm technology.

PEOPLE

RESEARCHERS

Krste Asanovic


VISITORS

Miquel Moretó Planas

STUDENTS AND INTERNS

Rimas Avizienis
Scott Beamer
Alex Bishara
Chris Celio
Henry Cook
Andrew Waterman
Richard Xia

INTERNATIONAL PROGRAMS





INTERNATIONAL PROGRAMS

ICSI is founded upon the principles of international partnership and cooperation. The Institute was established as a joint project of UC Berkeley and Germany's Research Center for Informational Technology GmbH in 1986. Since then, the Institute has hosted visitors from Spain, Switzerland, Finland, Italy, the U.K., the Philippines, Israel, China, Japan, Brazil, Belgium, Canada, and Holland. The Institute currently has formal international programs with Finland and Germany, but visitors travel from all over the world for research stays at ICSI funded by academic, industrial, or government sponsors.

The Institute hosts a variety of international visitors at different stages in their careers. Many visitors come to the Institute on post-doctoral fellowships. Graduate students, visitors from industry, and senior academic researchers also contribute to the world-class research accomplished here. Through its international program, the Institute has earned a reputation for having a diverse institutional culture.

In addition to bringing to Berkeley the best research from across the world, the Institute's international program sends back to sponsor countries the newest techniques and innovations in computer science. Visitors return to their home countries with a network of contacts from academia and industry in the Bay

Area, the United States, and other sponsor countries, and they return with an understanding of new and groundbreaking research.

Funds for the international visitor program are primarily provided by government agencies, but sometimes come through a combination of direct government funding, indirect funding through specific government research projects or institutions, support of industrial partners, or custom sponsorship organizations combining several funding sources. Agreements in place in 2012 included:

- Finnish support provided by the Finnish Funding Agency for Technology and Innovation (Tekes) through Aalto University and the Helsinki Institute for Information Technology.
- German support provided by a government-industry consortium, the German Academic Exchange Service (DAAD), and organized by a sponsoring association, the Förderverein.

Each year, ICSI hosts visitors from around the world. In the following section, we highlight some of the researchers who arrived in 2012.

July 25, 2012

Singapore joins ICSI's visitor program

Researchers from Nanyang Technical University (NTU) in Singapore will visit ICSI's Speech Group over the next two years through a research training program established by a recently signed agreement. Beginning September 1, NTU will send researchers from its Temasek Laboratories to visit ICSI in order to collaborate on speech research techniques. In the first year of the agreement, NTU researchers will work on multilingual speech recognition and keyword search, and in the second, on speech recognition robustness under difficult acoustic conditions.



VISITORS FROM GERMANY

German support is provided by the German Academic Exchange Service (DAAD) and organized by a sponsoring association, the Förderverein. ICSI was established as a collaboration between ICSI and the German Federal Laboratory for Computer Science (GMD). The first German-funded postdocs arrived in 1989. Since then, ICSI has hosted hundreds of German researchers at all career levels. Here we highlight a few of the postdocs whose DAAD-funded postdoctoral fellowships began in 2012.

MICHAEL ELBERFELD – ALGORITHMS

Michael began his PhD studies at the University of Lübeck by designing algorithms for computational problems that arise in molecular biology. Later he worked on the complexity theory of space-efficient and parallel computations. While time-efficient algorithms are fast, they often need huge amounts of memory for storing and manipulating data structures. Space-efficient algorithms trade time for space to save memory. His these focused on what problems can be solved using a logarithmic amount of space.



DANIEL GÖHRING – SPEECH

Daniel received his diploma and PhD from Humboldt University in Berlin. He focused on artificial intelligence; he and his group programmed robots for (and won three times at) RoboCup Soccer, an annual competition that pits teams of four robots against each other in soccer games. The robots are manufactured by Sony and programmed by teams of computer scientists around the world. The robots have low-resolution cameras with small fields of vision. Daniel programmed the robots to understand where they were on the field, along with their teammates and the ball, using constraint satisfaction methods, which uses less computing power than other methods. In April 2009, he started working at the German Aerospace Center, where he simulated fluctuations of street traffic patterns that result from changes in socio-geographic factors, such as the price of gas. He then worked at Freie University on autonomous, self-driving cars. In addition to producing code that, for example, kept the cars stable at high speeds, he was also the test driver of the cars.



CHRISTOF LENG – NETWORKING

Christof received his PhD from the Technical University of Darmstadt. For his thesis, he worked on BubbleStorm, a search system for peer-to-peer networks. BubbleStorm is a rendezvous search method, which allows users to search for data by using complex queries, like keyword search or even regular expressions. By contrast, distributed hash tables require that exact keys of values be known. In the simplest example of a rendezvous search, the nodes of a network are arranged in a grid and a piece of data is replicated on all points in a column. When a query is sent, it goes to all points along a row. The grid approach is not flexible enough to cope with large scale peer-to-peer networks, particularly as participants in the network come and go with high frequency. BubbleStorm deals with these complications by using stochastic algorithms. Christof identified different ways of replicating and updating data in this stochastic environment and used BubbleStorm for building an online game, a wiki, and a file-sharing program.



ERIK RODNER – VISION

Erik received his diploma and PhD at the University of Jena. His thesis work was on the use of machine learning for computer vision and specifically for category recognition. In category recognition, a computer system is asked to label a picture with its category. Erik was interested in reducing the amount of training data a system needs to do that. His system automatically took what it knew from other categories and applied it to the task of recognizing a new category. This allows the system to apply its knowledge about similar categories. He used Bayesian and kernel methods to accomplish this.





INTERNATIONAL PROGRAMS

DANIEL WARNEKE – NETWORKING

Daniel received his diploma from the University of Paderborn and went on to the doctoral program in computer science at the Berlin Institute of Technology. His thesis was on massively parallel data processing on Infrastructure as a Service platforms. His work examined and improved the efficiency of data-intensive applications on virtualized pay-as-you-go cloud systems like Amazon EC2, for example by exploiting the flexible cost model of those clouds for resource management. Daniel was a member of the Stratosphere research project. He wrote the parallel data processing framework Nephele, which now provides the foundation for the Stratosphere system, a rich open source software stack for next-generation big data analytics.



At ICSI in 2012:

BERNHARD AMANN – NETWORKING

Bernhard Amann received his diploma and PhD in computer science from the Technical University of Munich. For his thesis, he explored ways to provide different levels of access to files kept in decentralized file systems, groups of computers under the control of many different people. This can be complicated as the computers are not under unified control. Bernhard used algorithms typically used in other contexts – for example, in Blu-ray DVDs – to encrypt file keys and allow sharing of files across these computers in safe ways. He also worked on ways to ensure files are kept secure in the event that one or more of the computers in the system die, while not consuming too many computational resources. At ICSI, he will pursue his interest in network intrusion detection. He works on Bro, an open source network-monitoring framework developed at ICSI that helps defend networks against attacks.

OLIVER CULO – AI (FRAMENet)

Oliver Culo attended Saarland University, where he received his diploma in computational linguistics and his PhD in machine translation. His thesis work focused on developing ways of automatically comparing verb valence between English and German using parallel corpora. From a syntactic viewpoint, verb valence controls number, form, and function of phrases – for example, the subject, direct object, and indirect object. Oliver was interested in comparing the structures of English and German sentences and how verb valence may change when translating between the two languages. He's also generally interested in process-based research: his research group studied people as they read and translated texts on computers, tracking their eye movements and key strokes, and recording their general behavior on video. At ICSI, he worked with researchers in the FrameNet Project, who are building a lexical database based on frames – descriptions of events, relations, or entities and the participants in them.



GERARD DE MELO – AI (FRAMENet)

Gerard de Melo studied computer science with a minor in philosophy at Goethe University in Germany, also spending a year at the University of Montreal. He then joined the Max Planck Institute for Informatics and received his doctoral degree from Saarland University in Saarbrücken, Germany. His thesis work focused on ways of automatically creating large collections of knowledge about words in different languages. His database included descriptions of words in many languages and the relationships between them. One application of such work is to query expansion, a method of improving searches by including synonyms and related words (such as "automobile" for "car"). At ICSI, he works on the cross-lingual extension of FrameNet, a lexical database of words and how they are used. He also works on multi-lingual natural language processing tools.



MATTHIAS KIRCHNER – VISION

Matthias Kirchner works in image forensics, a field that tries to determine whether photos have been manipulated. The work has applications for both courtrooms, where photos may be used as evidence and so must be accurate, and newspapers, which may inadvertently print manipulated photos.

Matthias studied image forensics at the Technical University of Dresden, where he focused on geometric manipulation of photos. This occurs when a photo or part of a photo is resized or skewed. Matthias's PhD work, which was supported by a scholarship from Deutsche Telekom, used statistical analysis to find whether photos had been manipulated in this way. At ICSI, he focuses on local manipulation – that is, manipulation of certain areas of photos.



EMANUEL KITZELMANN – AI

Emanuel Kitzelmann did his PhD at the University of Bamberg in inductive synthesis to infer recursive functional programs. He built programs that could, given initial and final sets of information, infer what was done to the data in the meantime. He is interested in applying this sort of inductive reasoning to automated planning (for example, the Towers of Hanoi problem, in which disks of ascending sizes must be moved from one pile into another according to certain rules). He's generally interested in artificial intelligence and machine learning.



NILS PETERS – CAMPUS AFFILIATION (CNMAT)

While in Berkeley, Nils Peters collaborated with the Center for New Music and Audio Technologies (CNMAT) at UC Berkeley. Nils worked on his master's thesis in Graz, Austria, and received his PhD from McGill University in Montreal, Canada. His background is audio engineering with



an emphasis on room acoustics, signal processing, and human-computer interaction. Nils also worked as an audio engineer in recording and post-production and is also a co-developer of Jamoma, an open-source project for interactive media processing. His master's work focused on creating a wireless binaural headphones that tracks the user's head movements to dynamically adjust the presented audio signals so that the sound sources stay in place. His dissertation was about modeling and perception of spatialized sounds. While at ICSI and CNMAT, he worked on a microphone array system built by CNMAT and Meyer Sound. The array is approximately the size and shape of a soccer ball and is covered with 150 microphones. Nils hopes to develop real-time parallelization algorithms to process the data captured with the 150 microphones. These algorithms make it possible to estimate acoustical data about the room in which the array is placed -- e.g., to determine where sound sources and room reflections are coming from.

ANDREAS WUNDSAM – NETWORKING

Andreas Wundsam worked with Scott Shenker. He received his diploma at the Technical University of Munich and his PhD from the Technical University of Berlin in computer science. His thesis work focused on improving control and troubleshooting of operational networks. Many networks have grown extraordinarily complicated since they were first established, and for some, troubleshooting methods have not kept up. Andi worked to find network-centric and minimally invasive troubleshooting methods that could be implemented on real networks with a reasonable amount of effort, rather than methods that require the complete overhaul of a network. In particular, he has been using Virtual Network and OpenFlow as enabling technologies. For instance, with Mirror VNet, a production service inside of a VNet is paired with a Mirror VNet in identical state and configuration. This enables safer troubleshooting or upgrades to the system, as the maintenance work can be performed on the Mirror VNet in isolation from the production. In another scenario, he used the OpenFlow communication protocol to create a tool that allowed him to selectively record network traffic flow and later to replay the traffic in order to find causes of network problems.





INTERNATIONAL PROGRAMS

VISITORS FROM FINLAND

Finnish support is provided by the Finnish Funding Agency for Technology and Innovation through Aalto University and the Helsinki Institute for Information Technology.


SEPPO ENARVI – SPEECH

Seppo is a licentiate candidate at Aalto University, where he also received his master's degree. His master's thesis incorporated computer vision work he performed for the electronics company Genera. While at Genera, Seppo helped develop a system that scanned wood logs for defects, such as rot. After completing his master's, he continued working for Genera and taking classes in machine learning. Since last year, he has focused on speech recognition, developing dictation applications specifically for dentists. This involves both command dictation (dictation about status of a patient's teeth) and report dictation, which may include other information as well. He is generally interested in medical speech recognition, in addition to computer vision and machine learning.



KALLE PALOMÄKI – SPEECH

Kalle received his PhD from the Helsinki University of Technology (which later became Aalto University) in 2005. His thesis had two themes. In his initial work, he worked with magnetoencephalography (MEG). MEG measures the magnetic fields caused by electrical currents in the human brain.



Kalle and his colleagues measured study participants' neural reactions to simple noises, such as soundbursts and vowels. He was also interested in noise-robust speech recognition and computational auditory scene analysis. An example of this is the cocktail problem: picking a sound or kind of sound – say, a particular speaker – out of background noise. This tends to be easy for human listeners, but not for machines. Generally he's interested in what can be learned from human auditory systems. Recently, he's also become interested in working with general audio, not necessarily just speech. This could include music and city noises.

MIKKO KURIMO – SPEECH

Mikko is the leader of the Speech Group at the Information and Computer Science Department at Aalto University. He received his PhD at the Helsinki University of Technology (which later became part of Aalto), focusing on training methods for hidden Markov models for speech recognition. Since his postdoc at IDIAP, he has been a professor and chief research scientist at Aalto. In addition to IDIAP, he has visited and worked with several research centers: University of Colorado at Boulder, University of Edinburgh, SRI, Cambridge University, and Nagoya Institute of Technology. His group at Aalto consists of 12 researchers concentrating in several aspects of speech recognition. His work includes unsupervised morpheme-based language modeling and adaptation of acoustic and language models for various speakers and speaking styles.



July 27, 2012

Postdoc Nils Peters Wins Best Paper Award

DAAD-funded postdoctoral fellow Nils Peters won a best paper award at the Sound and Music Computing Conference for his work on the Spatial Sound Description Interchanges Format (SpatDIF). Peters's research in signal processing and spatial acoustics aims to analyze and semantically describe fields of sound. He works with an array of 150 microphones built by UC Berkeley's Center for New Music and Audio Technologies and Meyer Sound. The microphone array senses the sounding objects in a room environment, and the captured data are processed by algorithms to identify the room environment and to classify and detect the location of the sounding objects. SpatDIF organizes these sound field descriptions in a structured way. Peters was presented with the award on July 14.

OTHER VISITORS

In addition to researchers who visit ICSI through official visitor agreements, visitors travel from across the world to work on specific projects.

JIAO ZHANG – NETWORKING

Jiao is visiting ICSI's Networking Group. She received her bachelor's degree from the Beijing University of Posts and Telecommunications in 2008, and she is now a graduate student at Tsinghua University in Beijing, China. At Tsinghua University, she first conducted research into wireless sensor networks, especially on designing and evaluating energy-efficient routing and data aggregation protocols. Later, she focused on traffic



management in data centers. Data centers have many unique features, including multiple paths between servers, a large number of short messages, and special communication patterns in the data-intensive scalable computing systems, such as many-to-one and many-to-many communication patterns in Map Reduce, Dryad, Spark, and so on. These features cause some challenges, for example, TCP incast, large latency of the short messages, and TCP outcast. She has built an analytical model to analyze the causes of TCP incast and proposed some solutions to solve it. She is interested in designing a novel transmission protocol for data center networks to solve all the problems faced by TCP and other traffic management problems in data centers.



ICSI researchers gather to celebrate the release of Director Roberto Pieraccini's book, *The Voice in the Machine*

April 25, 2012

Roberto Pieraccini Visits Tekes in Finland

Director Roberto Pieraccini met with collaborators and partners involved in ICSI's Finnish visiting program April 23–25 in Finland. Since the program's establishment over 10 years ago, ICSI has hosted more than 50 Finnish visiting researchers funded through the Finnish Funding Agency for Technology and Innovation.

March 12, 2012

Board Member Honored for Service to Italian Research Center

ICSI Board of Trustees member Wolfgang Wahlster has received the Eagle of Saint Venceslao, the highest recognition given by the Autonomous Province of Trento, Italy. Wahlster was recognized for more than 15 years of service to IRST, now called the Bruno Kessler Foundation. IRST, located in Trento, is one of the most prestigious research centers in Italy.



INTERNATIONAL PROGRAMS

ICSI VISITORS IN 2012

NETWORKING

Amin Tootoonchian
 Andreas Wundsam (Germany)
 Bernhard Amann (Germany)
 Christof Leng (Germany)
 Daniel Warneke (Germany)
 Dorgival Guedes (Brazil)
 Eemil Lagerspetz (Finland)
 Haixin Duan
 Jason Croft
 Jiao Zhang
 Junaïd Khalid
 Lorenzo DeCarli
 Mohan Dhawan
 Oana Goga
 Renata Teixeira
 Sam Whitlock
 Sasu Tarkoma (Finland)
 Sheharbano Khattak
 Shuang Hao
 Soumya Basu

SPEECH

Arlo Faria (Brazil)
 Daniel Göhring (Germany)
 David Imseng (Switzerland)
 David Suendermann
 Frantisek Grezl
 Hai Do (Singapore)
 Kalle Palomaeki (Finland)
 Korbinian Riedhammer (Germany)
 Mikko Kurimo (Finland)
 Seppo Enarvi (Finland)

VISION

Erik Rodner (Germany)
 Lorenzo Riano
 Matthias Kirchner (Germany)
 Seppo Virtanen (Finland)
 Stefanie Jegelka
 Tim Althoff
 Tobias Baumgartner

AI

Emanuel Kitzelmann (Germany)
 Gerard de Melo (Germany)
 Hiroaki Sato
 Malte Schilling (Germany)
 Oliver Culo (Germany)
 Sergio Guadarrama

ALGORITHMS

Roded Sharan
 Itamar Eskin
 Farhad Hormozdiari
 Ron Shamir
 Michael Elberfeld (Germany)

ARCHITECTURE

Miquel Moretó Planas

CAMPUS AFFILIATION / OTHER

Nils Peters (Germany)



Prof. Wolfgang Wahlster and Director Roberto Pieraccini with a group of German visitors and DAAD alumni.

TECHNICAL CONTENT





TECHNICAL CONTENT

NETWORKING

Internet Security

The last decade has seen dramatic transformations in both the capability and use of malicious software. Among the scourges now routinely inflicted by Internet attackers are distributed denial-of-service attacks, widespread information theft, click fraud, spamming, and phishing. ICSI pursues research spanning many dimensions of this important problem space, ranging from technologies to detect attacks and analyze their aftermath to approaches for measuring, understanding, and disrupting attacker operations and goals.

ICSI's Bro technology (<http://bro.org>) provides a powerful, high-performance platform for analyzing network activity in rich semantic depth. Bro provides the user with a domain-specific scripting language used to define real-time analysis of network traffic in terms of events reflecting different elements of the traffic's activity. While we use Bro heavily to support different research projects, it also provides operational 24/7 monitoring at a number of large sites, some of which have deployed racks housing scores of servers in order to deeply analyze near-10Gbps traffic streams in real-time. Bro-related research includes: new detection and analysis algorithms; capabilities for integrating disparate information sources and exporting Bro analysis results to other systems; assessment of the fidelity and stability of the TLS Public Key Infrastructure that underlies modern Web security; compiler and hardware architectures for achieving very high performance; and database technology to support efficient archive and forensic searches over vast volumes of activity information.

In another dimension, one of the bedrocks of modern attacks concerns the ability of attackers to wield large numbers of compromised Internet systems ("bots") as a single coordinated instrument (a botnet). Indeed, it would be fair to say that the

emergence of modern cyber-crime is largely dependent on the economies of scale provided by botnets and the ease with which this technology can be distributed, specialized, bought, and sold. To this end, ICSI researchers work with colleagues to undertake botnet infiltration efforts, developing and deploying technology that allows systems under researcher control to act as bots and join into a botnet's command-and-control infrastructure, enabling monitoring of botnet operations and, in some particular (carefully controlled) situations, perturbation of the botnet's functioning. These experiments provide unprecedented visibility into the employment of botnet technology.

The work on botnets has recently led to research efforts focusing more deeply on cybercrime. With UC San Diego and George Mason University, ICSI recently formed CESR, the Center for Evidence-Based Security Research (<http://evidence-basedsecurity.org>). This interdisciplinary effort takes the view that, while security is a phenomenon mediated by the technical workings of computers and networks, it is ultimately a conflict driven by economic and social issues that merit a commensurate level of scrutiny. Begun in 2012 and initially supported by a five-year grant from the National Science Foundation, CESR has a primary goal of tackling key social and economic elements of security: how the motivations and interactions of attackers, defenders, and users shape the threats we face, how they evolve over time, and how they can best be addressed. Some of the initial work in this area has focused on assessing the business of cybercrime by infiltrating it as purported customers, making carefully considered purchases of goods and services available on underground marketplaces. We then in turn analyze both the purchasing process (what infrastructure is required to support it) and the results (what light the provided goods shed on the process used by the attacker to acquire/create them).

NETWORKING

Internet Analysis

ICSI researchers have a long history of pursuing empirical studies of how the Internet functions. Currently, the group undertakes a number of efforts in this regard, spanning a number of different measurement perspectives and network performance concerns.

First, the last mile, and, increasingly, the last meter often dominate what users perceive as the Internet: the performance provided by their home network, local network equipment, and Internet Service Provider. Failures, outages, and limitations may arise from a variety of sources, including incorrect NATs, broken firewalls, deliberate ISP manipulations, and over-buffered access equipment. Assessing these concerns can greatly benefit from observing from a user-centric vantage point: the perspective that users themselves have when they interact with the network. Our Netalyzr network troubleshooting service allows anyone with a Java-enabled Web browser to evaluate their Internet connectivity's current status, contributing in turn to a wide-ranging measurement effort studying the Internet's health (see <http://netalyzr.icsi.berkeley.edu>). To date, it has been run more than 800,000 times by more than half a million Internet systems, providing a wealth of data from an otherwise rarely explored perspective. We are now also pursuing the development of Fathom (<http://fathom.icsi.berkeley.edu>), a general approach for turning Web browsers into powerful measurement platforms by adding measurement capabilities to Javascript, the ubiquitous Web programming language.

A second area of focus concerns understanding the behavior of enterprise networks. Despite their widespread use and major roles in facilitating modern economic productivity, the networks used internally by modern enterprises have received at best fleeting examination in the research literature. This project seeks to enhance the understanding of enterprise networks by extensively studying a 2~TB collection of packet traces gathered from within an enterprise with tens of thousands of hosts. The efforts include assuring the data's quality via calibration (which turns out to be technically quite challenging), characterizing and modeling application usage within the enterprise, and

assessing the degree to which enterprise networks effectively deliver their full potential performance.

A third aspect of studying today's Internet centers on privacy: how third parties can collect, analyze, and employ information collected about individuals. Efforts in this area include both multimedia analysis and textual analysis to detect and categorize information leakage, as well as work on developing browser extensions that better protect privacy. The first area explores how a user's account names, images, tweets, and other posted content can enable third parties to de-anonymize users, as well as the privacy implications of multimedia image analysis. Our work on browser extensions demonstrates the viability of providing a browser cookie policy that prevents social media elements from tracking users without incurring any significant degradation to the functionality these elements provide.

A final dimension of Internet functioning that ICSI studies concerns censorship: the suppression or manipulation of Internet connectivity to control what information users can access. Around the world, Internet censorship is growing in scale, scope, and sophistication. National deployments affect millions of people, stifling the spread of ideas and the free flow of information. Aggressive censorship systems undermine the infrastructure and foundational protocols of the Internet, employing techniques such as DNS poisoning and TCP RST injection to block access to content, essentially launching "spoofing" attacks to interfere with otherwise legitimate communication. In mechanistic terms, government-controlled censorship systems arguably "attack" many more people than any other malicious actors. To this end, ICSI researchers study Internet censorship as practiced at scale, including its vulnerabilities. That is, our efforts seek to understand the workings of such systems when operationally deployed on very large networks, such as by nation-states, including the degree to which network users can leverage weaknesses in the functioning of the systems in order to directly evade the imposition of censorship.



TECHNICAL CONTENT

NETWORKING

Software-Defined Networking (SDN)

Software-defined networking is a new approach to network control planes that was invented at ICSI jointly with colleagues at Stanford and later Nicira. While a relatively recent development, with the first full-blown realization in 2008 comprising the NOX network operating system and the OpenFlow forwarding model, SDN has already gathered wide support from industry. The standards body for SDN—in the form of the Open Networking Foundation (ONF)—now has more than 70 member companies, including most of the major players in networking. SDN has also garnered broad interest from academia, with many workshops devoted to the topic and an increasing number of SDN submissions to major conferences.

Why the fuss over SDN? Networking is a notoriously slow moving field when it comes to basic infrastructure, so why is there such intense interest in SDN commercially and academically? The reason lies in the three separate ways SDN would change the status quo:

- Better modularity: For the first time, SDN would introduce layers into the control plane, much the same way layering has guided the development of the data plane. These basic layers are the network operating system (NOS), network management applications written on top of the NOS, and the forwarding model that provides the interface through which the NOS can control the physical switches and routers. This modularity enables greater code reuse, so new network functionality can be deployed through new network management applications, without having to rewrite the complicated distributed protocols used in the NOS nor the detailed forwarding behavior of the switches.
- New deployment models: SDN enables new functionality to be deployed through software running on top of the NOS. Moreover, acceptance of a common forwarding model (such as OpenFlow) means that this functionality can be deployed regardless of which brand of switches is being used. Thus, new functionality need not be tied to the providers of networking hardware, and can evolve at the speed of software.
- New abstractions for operators: Once this clean modularity is in place, one can build further abstractions on top of it to ease the job of network operators. One abstraction in particular, network virtualization, simplifies network management by allowing network operators to configure a logical network. The network virtualization solution then “compiles” these logical networks into configurations of the physical hardware.

ICSI has a broad and vibrant research program in SDN, covering areas such as how to scale SDN to global networks, how to troubleshoot SDN networks, how to formalize the consistency requirements of network control planes, and how one can improve the basic SDN architecture.

February 16, 2012

Scott Shenker Elected to National Academy of Engineers

Scott Shenker, leader of the Networking Group, has been elected as a member of the National Academy of Engineering, one of the highest distinctions accorded to engineers. Shenker was recognized for his contributions to Internet design and architecture. A professor at UC Berkeley, he is also a member of ICSI's Board of Trustees, a founding member of the ICSI Center for Internet Research, and a fellow of the ACM and the IEEE.

NETWORKING

Future Internet Architectures

Over the past decade there have been many attempts at clean-slate redesigns of the Internet architecture. These efforts have usually focused on addressing one or more of the current architecture's functional shortcomings, such as its poor security or its host-centric rather than content-centric design. These clean-slate activities, the most recent of which is NSF's Future Internet Architecture program, have produced, and will continue to produce, a wealth of insight into how one might build an Internet that better realizes these various features.

The recent architectural work at ICSI takes a radically different but wholly complementary approach, focusing not on specific functional improvements but on the broader issue of fostering architectural innovation. The reasoning behind this choice is simple: in most areas where the Internet architecture is known to be deficient, such as reliability and security, the literature is replete with proposals for how the architecture might be dramatically improved. However, almost none of these proposals has been realized in the current infrastructure because they face insurmountable deployment barriers—barriers that are largely due to the current Internet's lack of architectural modularity—and the literature on how to overcome this architectural impasse is comparatively barren. This suggests that the biggest intellectual challenge facing the current architecture is not a particular functional deficiency, but its inability to gracefully accommodate innovation.

Moreover, the need to support architectural innovation is of fundamental importance because any architecture, if it is unable to evolve, will eventually be found wanting in one respect or another. And yet, despite this need, we know more about how to build a more secure and reliable Internet than we do about how to build a more evolvable Internet. In short, while we know how to improve many of the Internet's features, we know far less about how to improve its overall future.

Our architectural efforts focus on architecting for innovation, with the goal of allowing the Internet to evolve over time and adapt to future requirements. The key design challenge is to identify the absolutely minimal portion of the design that needs to be universally agreed upon (and therefore fixed for long periods of time) and use this core as a foundational framework that allows the rest of the architecture to evolve more freely.

Our work here involves three major thrusts:

- How can host software be re-architected so that new architectures can be more easily used by hosts?
- How can the basic interdomain protocols be re-architected so that several can coexist simultaneously?
- How can SDN be leveraged to ease the deployment of new architectures?

The results of our investigations suggest that a highly evolvable Internet could be easily obtained with a few minor modifications.

TECHNICAL CONTENT


NETWORKING

Studying Twitter Spam's Use in Political Censorship

Spammers who posted almost half a million Twitter messages in order to silence debate over Russia's election in December 2011 likely purchased fraudulent accounts in bulk and posted the tweets from botnets, groups of malware-infected computers under the command of a single person. According to Networking researchers, the campaign took advantage of an underground economy based on spam, a phenomenon that researchers are studying in an attempt to improve methods of eliminating spam.

Within forty-eight hours of Russia's parliamentary election on December 4, 2011, thousands of people rallied to protest the voting fraud they claimed Russia's ruling party had committed; tens of thousands more protested through social media. Between December 5 and 6, Twitter users posted more than 800,000 messages containing hashtags – words or phrases added to tweets and aggregated in search results – related to the elections. Of those, nearly half were posted to accounts later identified by Twitter as fraudulent. The messages, many of which were garbled nonsense, diluted the legitimate tweets returned by searches for hashtags.

In April, Kurt Thomas of UC Berkeley and Chris Grier and Vern Paxson of ICSI's Networking Group presented an in-depth analysis of



the spam campaign at the USENIX Workshop on Large-Scale Exploits and Emergent Threats. They found that 99.5 percent of the suspended accounts that had posted tweets about the election were registered under mail.ru email addresses and followed certain naming patterns. Applying these patterns to all Twitter accounts under mail.ru addresses, the researchers found nearly a million accounts likely to be fraudulent, only 20 percent of which Twitter has suspended. The large number of spam accounts suggests that they were purchased from an online marketplace that registers and sells accounts in bulk.

In addition, the researchers also found that the IP addresses used to post spam about the election were far more dispersed around the globe than those used to post to legitimate messages, which tended to originate in Russia. The spam IP addresses were also more likely to appear on the Composite Blocking List, which comprises IP addresses flagged for sending spam and malware. These findings suggest that spammers sent their tweets from machines infected by malware and used to send spam email.

The infected machines form part of what researchers describe as an underground economy, in which, among other things, infected machines send spam email in order to sell illicit goods and infect other machines. Research into this economy may lead to improved ways to fight spam, as when Networking Group researchers and colleagues at UC San Diego identified three banks that authorize 95 percent of credit card sales of goods advertised by spam. Spam-based profits could be significantly reduced if credit card-issuing banks refused to settle transactions authorized by these banks.

The findings about the Russian election suggest that the monetization of spam can lead to a chilling effect on political conversation as well.

The good news? Twitter's default search returns messages ranked by, among other measures, their "relevance." While spam tweets may have drowned out legitimate political debate in real-time searches – which return messages in reverse chronological order – relevance searches returned 53 percent fewer spam tweets than real-time searches.

Related Paper: "Adapting Social Spam Infrastructure for Political Censorship." Kurt Thomas, Chris Grier, and Vern Paxson. Proceedings of the USENIX Workshop on Large-Scale Exploits and Emergent Threats (LEET), San Jose, California, April 2012. https://www.icsi.berkeley.edu/icsi/publication_details?n=3297

NETWORKING

Other Topics in Networking

ICSI maintains an active research program in more traditional networking topics such as routing, congestion control, and the like. Recent advances here include (i) the development of Data-Driven Connectivity, which enables the data plane to re-establish connectivity without waiting for the control plane to converge and (ii) a new congestion control algorithm for datacenters, called pFabric and developed jointly with researchers at Stanford, which provides near optimal performance while greatly simplifying the control algorithms and significantly reducing buffer sizes.

In addition, ICSI's Big Data Computing Paradigms effort, which will soon move to the newly created ICSI Initiatives Group, represents joint work with Berkeley's AMP Lab, of which Scott Shenker is a member. The work here involves designing new methods of organizing computation in the large clusters used to process Big Data. Areas of interest include new methods of resource allocation (e.g., Mesos), new programming abstractions (e.g., RDDs), realizations of those abstractions (e.g., Spark), higher-level functionality (e.g., Streaming, SQL queries), debugging infrastructure, and various performance issues (particularly as they relate to caching).

Publications

M. ALIZADEH, S. YANG, S. KATTI, N. McKEOWN, B. PRABHAKAR, AND S. SHENKER. Deconstructing Datacenter Packet Transport. Proceedings of the 11th ACM Workshop on Hot Topics in Networks (HotNets XI), Redmond, Washington, pp. 133-138, October 2012.

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.

G. ANANTHANARAYANAN, A. GHODSI, S. SHENKER, AND I. STOICA. Why Let Resources Idle? Aggressive Cloning of Jobs with Dolly. Proceedings of the 4th USENIX Workshop on Hot Topics in Cloud Computing (HotCloud '12), Boston, Massachusetts, June 2012.

G. ANANTHANARAYANAN, A. GHODSI, A. WANG, D. BORTHAKUR, S. KANDULA, S. SHENKER, AND I. STOICA. PACMan: Coordinated Memory Caching for Parallel Jobs. Proceedings of the 9th USENIX Symposium on Networked Systems Design and Implementation (NSDI '12), San Jose, California, April 2012.

M. CASADO, T. KOPONEN, S. SHENKER, AND A. TOOTOONCHIAN. Fabric: A Retrospective on Evolving SDN. Proceedings of the ACM SIGCOMM Workshop on Hot Topics in Software Defined Networking (HotSDN), Helsinki, Finland, August 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.

L. DiCICCIO, R. TEIXEIRA, M. MAY, AND C. KREIBICH. Probe and Pray: Using UPnP for Home Network Measurements. Proceedings of the 13th Passive and Active Measurement Conference (PAM 2012), pp. 96-105, Vienna, Austria, March 2012.

H. DUAN, N. WEAVER, Z. ZHAO, M. HU, J. LIANG, J. JIANG, K. LI, AND V. PAXSON. Hold-On: Protecting Against On-Path DNS Poisoning. Proceedings of the Conference on Securing and Trusting Internet Names (SATIN 2012), Teddington, United Kingdom, March 2012.

C. ENGLE, A. LUPHER, R. XIN, M. ZAHARIA, M. FRANKLIN, S. SHENKER, AND I. STOICA. Shark: Fast Data Analysis Using Coarse-grained Distributed Memory. Demo, ACM SIGMOD/PODS Conference, Scottsdale, Arizona, May 2012.

J. FEIGENBAUM, B. GODFREY, A. PANDA, M. SCHAPIRA, S. SHENKER, AND A. SINGLA. On the Resilience of Routing Tables. Proceedings of the 31st Annual ACM SIGACT-SIGOPS Symposium on Principles of Distributed Computing (PODC '12), Madeira, Portugal, July 2012.

J. FEIGENBAUM, B. GODFREY, A. PANDA, M. SCHAPIRA, S. SHENKER, AND A. SINGLA. Brief Announcement: On the Resilience of Routing Tables. Proceedings of the 31st Annual ACM SIGACT-SIGOPS Symposium on Principles of Distributed Computing (PODC '12), Madeira, Portugal, pp. 237-238, July 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.

D. GUPTA, A. SEGAL, A. PANDA, G. SEGEV, M. SCHAPIRA, J. FEIGENBAUM, J. REXFORD, AND S. SHENKER. A New Approach to Interdomain Routing Based on Secure Multi-Party Computation. Proceedings of the 11th ACM Workshop on Hot Topics in Networks (HotNets XI), Redmond, Washington, pp. 37-42, October 2012.

T. HALVORSON, J. SZURDI, G. MAIER, M. FELEGYHAZI, C. KREIBICH, N. WEAVER, K. LEVCHENKO, AND V. PAXSON. The BIZ Top-Level Domain: Ten Years Later. Proceedings of the 13th Passive and Active Measurement Conference, Vienna, Austria, March 2012.



TECHNICAL CONTENT

- A. JANIN, G. FRIEDLAND, AND R. SOMMER. Content based Privacy for Consumer Produced Multimedia. ICSI Technical Report TR-12-012, October 2012.
- D. LEVIN, A. WUNDSAM, B. HELLER, N. HANDIGOL, AND A. FELDMANN. Logically Centralized? State Distribution Tradeoffs in Software Defined Networks. Proceedings of the Workshop on Hot Topics in Software Defined Networks (HotSDN '12) at ACM SIGCOMM 2012, Helsinki, Finland, August 2012.
- L. MARTIGNONI, P. POOSANKAM, M. ZAHARIA, J. HAN, S. MCCAMANT, D. SONG, V. PAXSON, A. PERRIG, S. SHENKER, AND I. STOICA. Cloud Terminal: Secure Access to Sensitive Applications from Untrusted Systems. Proceedings of the USENIX Annual Technical Conference (ATC), Boston, MA, June 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.
- D. MCCOY, A. PITSILLIDIS, G. JORDAN, N. WEAVER, C. KREIBICH, B. KREBS, G. M. VOELKER, S. SAVAGE, AND K. LEVCHENKO. PharmaLeaks: Understanding the Business of Online Pharmaceutical Affiliate Programs. Proceedings of the 21st USENIX Security Symposium, Bellevue, Washington, August 2012.
- A. J. OLINER, A. P. IYER, E. LAGERSPETZ, I. STOICA, AND S. TARKOMA. Carat: Collaborative Energy Bug Detection. Poster and demo at the proceedings of the 9th USENIX Symposium on Networked Systems Design and Implementation (NSDI '12), San Jose, California, April 2012.
- B. RAGHAVAN, M. CASADO, T. KOPONEN, S. RATNASAMY, A. GHODSI, AND S. SHENKER. Software-Defined Internet Architecture: Decoupling Architecture from Infrastructure. Proceedings of the 11th ACM Workshop on Hot Topics in Networks (HotNets XI), Redmond, Washington, pp. 43-48, October 2012.
- B. RAGHAVAN, D. IRWIN, J. ALBRECHT, J. MA, AND A. STREED. An Intermittent Energy Internet Architecture. Proceedings of the Third International Conference on Future Energy Systems (e-Energy 2012), Madrid, Spain, May 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.
- C. ROSSOW, C. J. DIETRICH, C. KREIBICH, C. GRIER, V. PAXSON, N. POHLMANN, H. BOS, AND M. VAN STEEN. Prudent Practices for Designing Malware Experiments: Status Quo and Outlook. Proceedings of the 33rd IEEE Symposium on Security and Privacy (S&P 2012), San Francisco, California, pp. 65-79, May 2012.
- S. SHENKER. Software-Defined Networking: History, Hype, and Hope. Proceedings of the Conference on Large-Scale Distributed Systems and Middleware (LADIS 2012), Madeira, Portugal, July 2012. Also in the proceedings of the Brazilian Symposium on Computer Networks and Distributed Systems (SBRC 2012), Ouro Preto, Brazil, April 2012.
- C. SHUE, A. KALAFUT, M. ALLMAN, AND C. TAYLOR. On Building Inexpensive Network Capabilities. ACM SIGCOMM Computer Communications Review, Vol. 42, No. 2, pp. 72-79, April 2012.
- K. THOMAS, C. GRIER, AND V. PAXSON. Adapting Social Spam Infrastructure for Political Censorship. Proceedings of the Fifth USENIX Workshop on Large-Scale Exploits and Emergent Threats (LEET '12), San Jose, California, April 2012.
- A. TOOTOONCHIAN, S. GORBUNOV, Y. GANJALI, M. CASADO, AND R. SHERWOOD. On Controller Performance in Software-Defined Networks. Proceedings of the USENIX NSDI Workshop on Hot Topics in Management of Internet, Cloud, and Enterprise Networks and Services (Hot-ICE), San Jose, California, April 2012.
- A. WUNDSAM AND M. YU. NOSIX: A Portable Switch Interface for the Network Operating System. ICSI Technical Report TR-12-013, October 2012.
- R. XIN, J. ROSEN, M. ZAHARIA, M. J. FRANKLIN, S. SHENKER, AND I. STOICA. Shark: SQL and Rich Analytics at Scale. Technical Report, UCB/EECS-2012-214, University of California at Berkeley, Department of Electrical Engineering and Computer Science, arXiv:1211.6176 [cs.DB], November 2012.
- M. ZAHARIA, M. CHOWDHURY, T. DAS, A. DAVE, J. MA, M. MCCAULEY, M. J. FRANKLIN, S. SHENKER, AND I. STOICA. Resilient Distributed Datasets: A Fault-Tolerant Abstraction for In-Memory Cluster Computing. Proceedings of the 9th USENIX Symposium on Networked Systems Design and Implementation (NSDI 2012), San Jose, CA, April 2012.
- M. ZAHARIA, M. CHOWDHURY, T. DAS, A. DAVE, J. MA, M. MCCAULEY, M. J. FRANKLIN, S. SHENKER, AND I. STOICA. Fast and Interactive Analytics Over Hadoop Data with Spark. USENIX ;login;, Vol. 34, No. 4, pp. 45-51, August 2012.
- M. ZAHARIA, T. DAS, H. LI, S. SHENKER AND I. STOICA. Discretized Streams: An Efficient and Fault-Tolerant Model for Stream Processing on Large Clusters. In the Proceedings of the 4th USENIX Conference on Hot Topics in Cloud Computing (HotCloud '12), Boston, MA, June 2012.
- M. ZAHARIA, S. KATTI, C. GRIER, V. PAXSON, S. SHENKER, I. STOICA AND D. SONG. Hypervisors as a Foothold for Personal Computer Security: An Agenda for the Research Community. EECS Department, UC Berkeley Technical Report No. UCB/EECS-2012-12, Berkeley, California, January 2012.

SPEECH

Speech Recognition and Related Topics

PITCH DETECTION

We have developed a trainable pitch detector that is based on classification of subband autocorrelation features (Fig. 1), now called the subband autocorrelation classification (SACc) pitch tracker [1]. By training a single large multi-layer perceptron (MLP) on the normalized, PCA-reduced autocorrelations of 48 sub-bands, we created a pitch tracker that significantly outperforms a state-of-the-art pitch tracker, particularly for band-limited, noisy speech. This has proved useful for a number of the speech processing tasks, including those described below.

SPEECH ACTIVITY DETECTION (SAD)

Using the pitch detector described above, as well as the MLP-processed Gabor-filtered features described in previous years' reports, we have generated features for our SRI partners that have significantly improved their SAD systems. We also analyzed the nature of features that appeared to be most effective for SAD in a study presented at Interspeech [2]. In related work we developed an approach to effectively using the allowable level of false alarms and misses as resources in choosing the best feature for each speech segment [13].

SPEAKER DIARIZATION

Speaker diarization is the task of partitioning an input stream into speaker homogeneous regions, or in other words, to determine "who spoke when." ICSI contributed to speaker diarization research for many years and the approach manifested itself among the state-of-the-art [3]. This year, we continued our work on comparing different uni-modal diarization approaches, namely all the ones participating in NIST RT'09 and one for IDIAP to various multimodal approaches. We published some of the results of this study [4] and the extensive discussion will be part of Mary Knox's PhD thesis.

SPEAKER RECOGNITION

In 2012, ICSI participated in the biennial international evaluation of speaker recognition technologies conducted by the National Institute of Standards and Technology (NIST). The focus of our work was to be able to handle large amounts of data, improving the system's robustness through deploying parallelism while

using off-the-shelf tools to create a state-of-the-art i-vector speaker recognition system. Other work has focused on gaining a better understanding of the effects of phonetic structure in speaker recognition models.

SPEECH RECOGNITION

We developed a hybrid of CMU's Power Normalized Cepstral Coefficients (PNCC) with our own Gabor filter features to create a speech recognition features set that was more robust to additive noise than either approach [12]. In particular, we used the power normalized spectra from PNCC as the raw materials for spectro-temporal filtering by Gabor filters (Fig. 2) followed by transformation by a discriminatively trained MLP. Also, as noted above, we made significant progress in diagnosing the deficiencies of the acoustic models in HMM-based speech recognition [11]. As an additional component of this diagnostic study, we conducted a survey of the automatic speech recognition (ASR) community and of the relevant literature as part of an effort to catalog other problems with current models as a first step toward developing the ASR system of the future. We also began to explore the relationship between the recent "deep learning approaches" to speech recognition and prior efforts incorporating neural networks, as summarized in [5]. Finally, as part of our continuing exploration of biologically-inspired methods for ASR signal processing, we co-authored a survey with a CMU colleague on such approaches [6].

RAPID DEVELOPMENT AND PORTING OF ASR SYSTEMS FOR NEW LANGUAGES

We have developed a modular Python-based "System for Running Systems" (SRS), which helps to run, debug, restart, and archive complex experiments, exploits compute cluster infrastructure that has hooks for data storage (e.g., local vs. networked; temp vs. archival), and uses configuration files to control interaction between steps, for parameter settings. We have built a complete ASR training and recognition system using HTK as the core recognition software, but also using many other modules, such as neural network training and evaluation. The system takes a standard IARPA language pack for any language, which includes audio, transcriptions, and a lexicon, and runs on it to produce a recognition system.



TECHNICAL CONTENT

SPEECH

Multimedia and General Audio

MULTIMODAL LOCATION ESTIMATION

This project is a collaboration with UC Berkeley faculty member Kannan Ramchandran. Location estimation is the task of estimating where digital media was recorded. Our project aims to use the GPS-tagged media available on the Web as a training set for an automatic location estimator. We continued our development of the ICSI location estimation system, which participated in the European MediaEval 2012 Placing task. The new system is better able to cope with sparse training data. The algorithm was published at ICME 2012 [7]. Results are illustrated in Fig. 3.

VIDEO CONCEPT DETECTION

Given the exponential growth of videos published on the Internet, mechanisms for clustering, searching, and browsing large numbers of videos have become a major research area. More importantly, there is a demand for event detectors that go beyond simply finding objects to detecting more abstract concepts, such as "feeding an animal" or a "wedding ceremony." ICSI has been fortunate to become part of IARPA's Aladdin program. The program aims at describing the content of a "found" video based on a set of example videos. The computer learns concepts from example videos and then recounts the concepts seen in the query videos. The task is performed on a set of 100,000 "wild" Internet videos – videos freely available on the Web and unregulated for quality, size, or content. ICSI is working in a team together with SRI/Sarnoff, Carnegie Mellon University, University of Massachusetts, University of Central Florida, and Cycorp. ICSI's contributions are acoustic, visual, and multimodal. In 2012 we improved upon the 2011 approach by using the specific distribution properties of unsupervised audio clusters, which we called percepts [8].

PRIVACY IMPLICATIONS OF MULTIMEDIA RETRIEVAL

We are exploring the ways in which it is possible to aggregate public and seemingly innocuous information from different media and Web sites to attack users' privacy. The project seeks to help users understand the privacy implications of the information they share publicly on the Internet and to help them understand what control they can exercise over it. User-supplied content – in the form of photos, videos, and text – is a crucial ingredient to many Web sites and services today. However, many users who provide content do not realize that their uploads may be leaking personal information in forms hard to intuitively grasp. Correlation of seemingly innocuous information can create inference chains that tell much more about individuals than they are aware of revealing. We contend that adversaries can systematically exploit such relationships by correlating information from different sources in what we term global inference attacks: assembling a comprehensive understanding from individual pieces found at a variety of locations. Not only are such attacks already technically viable given the capabilities of today's multimedia content analysis and correlation technologies, but we also find business models that provide adversaries with powerful incentives for pursuing them. In collaboration with the Networking Group, we are therefore exploring the risks of multimedia retrieval to online privacy and also investigating possible countermeasures.

In 2012 we started research on matching anonymous accounts based on innocuous activity. We investigated chains with up to three links between Yelp, Twitter, and Flickr. The results of a study on more than 1 million records have been submitted and accepted to WWW2013 [9].

SPEECH

Multimedia and General Audio, cont.

TEACHING PRIVACY TO K-12


Based on NSF funding, we started an educational project in collaboration with Dan Garcia, a UC Berkeley (lecturer, that would teach the findings of our privacy research to K-12 students. Teenagers are the most intensive users of social media and yet the least aware of privacy issues. A curriculum is being developed in close collaboration with BFOIT. Initially, we developed 10 principles for social media privacy, which we are currently refining using expert knowledge from the community. We are developing a project Web site: <http://teachingprivacy.icsi.berkeley.edu>.

CROWDSOURCING INTELLIGENCE

Usually crowdsourcing is used for tasks that are easy to perform for humans but difficult for machines. In this NSF-sponsored work, we wanted to know if one could crowdsource tasks that are difficult to perform for humans and machines. We developed methods for finding Mechanical Turk participants for the manual annotation of the geo-location of random videos from the Web. This is a task that is hard for both humans as well as computers since it requires finding cues that indicate location -- which may or may not be in the video. Furthermore, we required high-quality annotations for this project, as we were attempting to establish a human baseline for future comparison with machine systems. Our idea for solving this problem is to qualify and select the crowdsourced workers. This is different from previous crowdsourcing paradigms that usually assumed the human crowd to be large set of dumb "mechanical" units. We worked on the varied difficulties we encountered while qualifying annotators and the steps that we took to select the individuals most likely to do well at our annotation task in the future. Our publication at ACM Multimedia [10] was very well received and resulted in invitations for publications in other places (see highlights).

PARALLELIZATION OF AUDIO ANALYTICS

In this new era dominated by consumer-produced media there is a high demand for Web-scalable solutions to multimedia content analysis. A compelling approach to making applications scalable is to explicitly map their computation onto parallel platforms. However, developing efficient parallel implementations and fully utilizing the available resources remains a challenge due to the increased code complexity, limited portability, and required low-level knowledge of the underlying hardware. Within our component of UC Berkeley's Par Lab, we developed PyCASP, a Python-based framework that automatically maps computation onto parallel platforms from Python application code to a variety of parallel platforms. PyCASP is designed using a systematic, pattern-oriented approach to offer a single software development environment for multimedia content analysis applications. Using PyCASP, applications can be prototyped in a few hundred lines of Python code and automatically scale to modern parallel processors. Applications written with PyCASP are portable to a variety of parallel platforms and efficiently scale from a single desktop graphics processing unit (GPU) to an entire cluster with little change to application code. We created a demo of a content-based music recommendation system based on the Million Song Dataset using Python and started to use the framework for video concept detection.



Dr. Gerald Friedland and his audio and multimedia team



TECHNICAL CONTENT

SPEECH

The Room Within the Sound

Researchers have developed a method to identify rooms using audio recordings that were captured in them. Nils Peters, a DAAD-funded postdoctoral fellow, and Speech researchers Gerald Friedland and Howard Lei used audio recordings from seven different rooms – a bedroom, library, studio, two churches, great hall, and classroom – to develop an acoustic profile for each. These profiles were based on audio features that are frequently used by speech systems to automatically recognize words.

The audio features of new unknown audio recordings, containing either speech or music, were then compared against these acoustic profiles. When there was no overlap between the training and test data, the system accurately identified the type of room 61 percent of the time when music was used and 85 percent when speech signals were used. By contrast, random chance would identify rooms with about 14 percent accuracy.

The work may help improve location estimation, which tries to automatically estimate the origins of videos, in situations where other techniques fail. GPS, for example, often fails indoors. Other applications include smart hearing aids able to automatically adapt to the acoustic environment.

The paper is related to ICSI's ongoing research that attempts to use multiple modalities, such as audio, text, and visual cues, to understand multimedia content automatically.

Related Paper: "Name That Room: Room Identification Using Acoustic Features in a Recording." Nils Peters, Howard Lei, and Gerald Friedland. Proceedings of ACM Multimedia 2012, Nara, Japan, October 2012. Available online at http://www.icsi.berkeley.edu/icsi/publication_details?n=3318.

2012 Highlights

OUCH PROJECT

Quantitative characterization of speech recognition errors. A large number of the errors in modern speech recognizers are due to the inaccuracy of the conditional independence assumptions in HMMs. Other sources of errors are the mismatch between training and test conditions and the signal representation. This study produces tools that are useful for future advances in ASR by providing more detailed diagnostics beyond simple word error rate.

SWORDFISH PROJECT

Speech recognition in under-represented languages. Development of tools for the generation of "language packs" for any written language, without large amounts of data. For more information about this project, see page 14.

July 1, 2012

Director Roberto Pieraccini Named 'Speech Luminary' by Speech Tech Magazine

Director Roberto Pieraccini has been named a 2012 speech luminary by Speech Technology magazine. The speech luminary awards are given each year to honor those whose creativity and drive have significantly influenced the speech technology industry.

SPEECH

Selected References

- B.-S. LEE AND D. ELLIS. Noise Robust Pitch Tracking by Subband Autocorrelation Classification. Proceedings of the 13th Annual Conference of the International Speech Communication Association (InterSpeech 2012), paper P3b.05, Portland, Oregon, September 2012.
- T.J. 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.
- X. ANGUERA, S. BOZONNET, N. EVANS, C. FREDOUILLE, G. FRIEDLAND, AND O. VINYALS. Speaker diarization: A review of recent research. *IEEE Transactions on Audio, Speech, and Language*, Vol. 20, Issue 2, pp. 356–370, January 2012.
- M. 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.
- N. MORGAN. Deep and Wide: Multiple Layers in Automatic Speech Recognition. *IEEE Transactions on Audio, Speech, and Language*, Vol. 20, Issue 1, pp.4–6, January 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.
- J. CHOI, G. FRIEDLAND, V. EKAMBARAM, AND K. RAMCHANDRAN. Multimodal Location Estimation of Consumer Media – Dealing with Sparse Training Data. Proceedings of the IEEE International Conference on Multimedia and Expo, Melbourne, Australia, pp. 43–48, July 2012.
- 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.
- O. GOGA, H. LEI, S. H. K. PARTHASARATHI, G. FRIEDLAND, R. SOMMER, AND R. TEIXEIRA. Exploiting Innocuous Activity for Correlating Users Across Sites. Proceedings of the World Wide Web Conference (WWW), Rio de Janeiro, Brazil, May 2013.
- 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.
- S. PARTHASARATHI, S.-Y. CHANG, J. COHEN, N. MORGAN, AND S. WEGMANN. The Blame Game in Meeting Room ASR: An Analysis of Feature vs Model Errors in Noisy and Mismatched Conditions. Accepted for ICASSP 2013, May 2013.
- S.-Y. CHANG, B. MEYER, N. MORGAN. Spectro-temporal features for noise-robust speech recognition using power-law nonlinearity and power bias subtraction. Proceedings of the IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP 2013), Vancouver, Canada, May 2013.
- T.J. TSAI AND N. MORGAN. Speech activity detection: an economics approach. Proceedings of the IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP 2013), Vancouver, Canada, May 2013.

Publications

- X. ANGUERA, S. BOZONNET, N. EVANS, C. FREDOUILLE, G. FRIEDLAND, AND O. VINYALS. Speaker Diarization: A Review of Recent Research. *IEEE Transactions on Audio, Speech, and Language Processing*, Vol. 20, Issue 2, pp. 356–370, February 2012
- A. BLACK, H. T. BUNNELL, Y. DOU, P. KUMAR, F. METZE, D. PERRY, T. POLZEHL, K. PRAHALLAD, S. STEIDL, AND C. VAUG. Articulatory Features for Expressive Speech Synthesis. Proceedings of the IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP 2012), Kyoto, Japan, March 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
- J. CHOI, G. FRIEDLAND, V. EKAMBARAM, AND K. RAMCHANDRAN. Multimodal Location Estimation of Consumer Media – Dealing with Sparse Training Data. Proceedings of the IEEE International Conference on Multimedia and Expo, Melbourne, Australia, pp. 43–48, July 2012
- 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
- G. FRIEDLAND. Cybercasing the Joint: Language Technologies, Multimedia Retrieval, and Online Privacy. Presented at the Language Technologies Institute Colloquium, Carnegie Mellon University, Pittsburgh, Pennsylvania, April 13 2012
- G. FRIEDLAND, J. FISCUS, T. HAIN, AND S. FURUI (EDS). Special Section on New Frontiers in Rich Transcription. *IEEE Transactions in Audio, Speech, and Language Processing*, Vol. 20, No. 2, February 2012
- G. FRIEDLAND, A. JANIN, AND L. GOTTLIEB. Narrative Theme Navigation for Sitcoms Supported by Fan-Generated Scripts. To appear in *Multimedia Tools and Applications*, Springer, 2012

TECHNICAL CONTENT

G. FRIEDLAND, A. JANIN, D. IMSENG, X. ANGUERA, L. GOTTLIEB, M. HUIJBREGTS, M. KNOX, AND O. VINYALS. The ICSI RT-09 Speaker Diarization System. *IEEE Transactions on Audio, Speech, and Language Processing*, Vol. 20, Issue 2, pp. 371–381, February 2012

G. FRIEDLAND AND F. VALENTE. Speaker Diarization. In *Multimodal Signal Processing: Human Interactions in Meetings*, S. Reynals, H. Bourlard, J. Carletta, and A. Popescu-Belis, eds., Cambridge University Press, June 2012

D. GILICK, S. WEGMANN, AND L. GILICK. Discriminative Training for Speech Recognition is Compensating for Statistical Dependence on the HMM Framework. *Proceedings of the IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP 2012)*, Kyoto, Japan, March 2012

D. GOEHRING, D. LATOTZKY, M. WANG, AND R. ROJAS. Semi-Autonomous Car Control Using Brain Computer Interfaces. *Proceedings of the 12th International Conference of Intelligent Autonomous Systems (IAS)*, Juju Island, Korea, June 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

P.-S. HUANG, R. MERTENS, A. DIVAKARAN, G. FRIEDLAND, AND M. HASEGAWA-JOHNS. How to Put It Into Words – Using Random Forests to Extract Symbol Level Descriptions from Audio Content for Concept Detection. *Proceedings of the IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP 2012)*, Kyoto, Japan, March 2012

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

H. LEI, J. CHOI, AND G. FRIEDLAND. Multimodal City-Verification on Flickr Videos Using Acoustic and Textual Features. *Proceedings of the IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP 2012)*, Kyoto, Japan, March 2012

H. LEI, B. T. MEYER, AND N. MIRGHAFORI. Spectro-Temporal Gabor Features for Speaker Recognition. *Proceedings of the IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP 2012)*, Kyoto, Japan, March 2012

X. LIU, J. L. HIERONYMUS, M. J. F. GALES, AND P. C. WOODLAND. Syllable Models for Mandarin Speech Recognition: Exploiting Character Language Models. In submission, 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

N. MORGAN. Deep and Wide: Multiple Layers in Automatic Speech Recognition. *IEEE Transactions on Audio, Speech, and Language Processing*, Vol. 20, Issue 1, pp. 7–13, January 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

N. PETERS, T. LOSSIUS, AND T. PLACE. An Automated Testing Suite for Computer Music Environments. *Proceedings of the 9th Sound and Music Computing Conference (SMC 2012)*, Copenhagen, Denmark, July 2012

N. PETERS AND S. McADAMS. A Perceptual Analysis of Off-Center Sound Degradation in Surround-Sound Reproduction Based on Geometrical Properties. *Journal of the Acoustical Society of America*, Vol. 131, No. 4, pp. 3256–3256. Also in the proceedings of Acoustics 2012, Hong Kong, May 2012

N. PETERS, J. SCHACHER, AND T. LOSSIUS. SpatDIF: Principles, Specification, and Examples. *Proceedings of the 9th Sound and Music Computing Conference (SMC 2012)*, Copenhagen, Denmark, July 2012

R. PIERACCINI. From AUDREY to Siri: Is Speech Recognition A Solved Problem? Presented at the Mobile Voice Conference, San Francisco, California, March 2012

S. RAVURI AND N. MORGAN. Easy Does It: Robust Spectro-Temporal Many-Stream ASR Without Fine Tuning Streams. *Proceedings of the IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP 2012)*, Kyoto, Japan, March 2012

R. M. STERN AND N. MORGAN. Features Based on Auditory Physiology and Perception. In *Techniques for Noise Robustness in Automatic Speech Recognition*, T. Virtanen, B. Raj, and R. Singh, Wiley Publishing, 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

T.J. 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

D. YU, G. HINTON, N. MORGAN, J.-T. CHIEN, AND S. SAGAYAMA. Introduction to the Special Section on Deep Learning for Speech and Language Processing. *IEEE Transactions on Audio, Speech, and Language Processing*, Vol. 20, Issue 1, pp. 4–6, January 2012

AI

MetaNet: A Multilingual Metaphor Extraction, Representation, and Reasoning System

The MetaNet project spans six institutions (ICSI, UC Berkeley, University of Southern California, UC San Diego, Stanford, UC Merced) and five disciplines (computer science, cognitive science, linguistics, psychology, and neuroscience) and currently has more than 35 students, postdocs, researchers, and faculty. In 2012, the team won a highly competitive grant and started work on a five-year project to build a computer system capable of understanding metaphors used in English, Persian, Russian, and Spanish.

MetaNet relies on the long and productive history within the ICSI AI Group of combining computational modeling with cognitive linguistic theories of metaphor. Over the last two decades, researchers from the Neural Theory of Language (NTL) project, including Professors Jerry Feldman, George Lakoff, and Srin Narayanan and several students, have combined biologically based computational models with cognitive linguistic analyses and experimental techniques (both behavioral and imaging) to demonstrate the ubiquity of metaphor, its connection to fundamental cognitive processes, and its use in everyday reasoning and in specialized discourse.



Prof. Jerome Feldman, a collaborator of the MetaNet project

The goal of IARPA's Metaphor Program is to build a system that extracts linguistic manifestations of metaphor (words and phrases that are based on metaphor) from text and interprets them automatically in four different languages. Researchers in ICSI's MetaNet project are doing this by building a multi-lingual metaphor repository that represents the network of conceptual metaphors and includes links to linguistic realizations. Users will be able to browse, navigate, annotate, and modify the repository, which will also provide programmatic access for metaphor extraction, analysis, and inference. The work is being done in American English, Iranian Persian, Russian as spoken in Russia, and Mexican Spanish.

In the first year of the project, the ICSI team developed the design and first implementation of a scalable methodology. Specific accomplishments of the first year include the following.

- 1.) A semi-supervised machine learning-based natural language processing framework that can extract metaphors automatically from large textual corpora in multiple languages with high precision ($>.7$). The system was able to extract more than 5,000 linguistic metaphors in the four languages of interest. As far as we are aware, this is the first multilingual metaphor extraction system and the numbers were very encouraging. The algorithm uses a combination of spectral clustering over dependency parses over large corpora and example seed annotations to extract the linguistic metaphors. The team also built a hierarchical factorized graph clustering algorithm that uses an unsupervised framework. The new algorithm was tested in English and the performance is comparable to the semi-supervised approach. Results are reported in a paper by ICSI MetaNet researcher Dr. Ekaterina Shutova, to be presented at NAACL 2013.

TECHNICAL CONTENT

2.) A fully functional repository with a hierarchical relational network of conceptual metaphors that includes more than 1,000 conceptual mappings and their associated frames, metaphors, role bindings, and relations. The conceptual metaphors span basic embodied metaphors that are hypothesized to be cross-linguistic as well as language- and culture-specific mappings related to the first-year target domain of Governance, set by IARPA. The repository is available in a variety of formats, including an RDF/OWL triple store and in MySQL. The team also developed a seamless interface for editing and annotation based on extensions to the Semantic MediaWiki. The repository is programmatically accessible through a JAVA API. The team built a visualization tool that can be used as a standalone JAVA application or as a plugin to Protégé, a standard open source ontology editing tool. Results are reported in several papers at the International Cognitive Linguistics Conference.



Prof. George Lakoff, a collaborator of the MetaNet project



Team members of the MetaNet project

3.) Analysis of the affective and inferential effects of metaphoric language. The strong connection between metaphor and bodily experience has implications for political and social discourse. A group at USC led by Professor Lisa-Zadeh collaborated with ICSI to design and conduct an fMRI imaging study on the ability of metaphoric discourse on governance to evoke strong emotional reactions such as fear and disgust. A group at UCSD led by Professor Ben Bergen used MEG with the same stimuli to provide a contrasting technique to correlate results. Initial results on both these experiments are encouraging, and if the analysis holds up this should result in major publications in the coming years on the role of metaphoric affect in decision-making. A Stanford team led by Professor Lera Boroditsky working with student and postdoc Paul Thibodeau, showed how different metaphors can lead people to completely different attitudes and decisions concerning major policy issues like crime. This work was reported in PLoS One in 2012.

The MetaNet project is supported by the Intelligence Advanced Research Projects Activity (IARPA) via Department of Defense US Army Research Laboratory contract number W911NF-12-C-0022.

AI

Digital Functioning in Underserved Populations

Researchers in the AI Group who are interested in designing for the disadvantaged are asking the public policy-related question "What is the best possible way to provide underserved populations with the necessary knowledge to use computers and the Internet in a way that is valuable and beneficial for them?" A new project builds on previous work, particularly the Hesperian and Templeton projects, on solutions to social problems that emerge with technology innovations.

In 2012, the group, led by Dr. Blanca Gordo, developed the sustainable broadband adoption theoretical framework to create a baseline for evaluating the Foundation for California Community Colleges (FCCC) statewide community program, California Connects (CC). CC is one of 44 Sustainable Broadband Adoption Programs (SBAs) funded through the American Recovery Reinvestment Act (ARRA) Title VI Broadband Technology Opportunities Program (BTOP) and administered by the U.S.

Department of Commerce's National Telecommunications and Information Administration (NTIA). This evaluation employs a cross-disciplinary conceptual framework, applying results from recent cognitive science research on framing and metaphor within the ICSI AI Group along with other results in social science that illustrate the interrelation of technology, innovations, and society. From this viewpoint, the effects of using technology are contextualized within an institutional and social structure and defined in terms of what technology makes possible. In other words, benefit is tied to the ways technology facilitates functioning in everyday life, saving users time and resources. The team examines the structure and effectiveness of CC in the context of California's social situation and needs of the target population. The research is ongoing, but the main results of the evaluation and recommendations for future use, will be released in 2013 to the FCCC and NTIA.

FrameNet

The FrameNet project is continuing to build a knowledge base of core English vocabulary based on the principles of frame semantics derived from the work of Professor Charles J. Fillmore. This lexical knowledge base is supported by more than 190,000 manually annotated examples from corpora of contemporary English. Similar projects are underway for more than eight other languages, and researchers around the world are using the freely available FrameNet data for a variety of natural language processing tasks.

The FrameNet team is collaborating with the MetaNet team to define frames and frame elements that can be used as source domains for metaphorical mappings. The connection between frames and metaphors is being further explored in a workshop organized by FrameNet group member Miriam Petruck, at the International Cognitive Linguistics Conference in 2013.

FrameNet also continued its collaboration with Decisive Analytics Corporation, a defense contractor in the Washington D.C. area. DAC builds systems for analyzing texts about military operations, and FrameNet has been creating new semantic frames and lexical units for this domain. This collaboration has

been funded by a series of SBIR grants to DAC, with FrameNet as a subcontractor. We were also able to carry out fairly extensive data integrity checking and error correction as part of this work.

In 2012, FrameNet members Miriam R. L. Petruck and Michael Ellsworth taught a 45-hour course at the International School on Frame Semantics and Technological Applications co-hosted by the Graduate Program in Linguistics of the Universidade Federal de Juiz de Fora and FrameNet Brasil. A total of 140 students participated, either in person or virtually, in the course, which focused primarily on FrameNet methodology and annotation practice.

FrameNet received two major awards from the community in 2012. Professor Fillmore won the ACL Lifetime Achievement award for his contributions to Frame Semantics and to the field of computational linguistics. Professor Fillmore and Dr. Collin Baker, the FrameNet project manager, won the prestigious Antonio Zampolli Prize, awarded by the European Language Resources Association for their contributions to language resources and language technology.

TECHNICAL CONTENT

AI

Meaningful Natural Language-Based Human-Computer Interaction

Funded by the Templeton Foundation and by the Office of Naval Research, a group led by Professors Jerry Feldman and Srinu Narayanan continued work on building systems capable of natural language interaction with computers. The project currently employs two graduate students, one software engineer, and two postdocs funded by the DAAD through ICSL's German visiting agreement. In 2012, the group had the following accomplishments.


1. Expressive language analysis (embodied construction grammar) and inference (coordinated probabilistic relational models) framework combining cognitive linguistics, construction grammar, first order probabilistic inference, and coordinated (concurrent and discrete/continuous state) event and process models.
2. First natural language analysis algorithm for deep semantic analysis and inference with
 - a. Event frames
 - b. Scales and polarity
 - c. Causatives and counterfactuals
 - d. Linguistic aspect
 - e. Figurative language (partly funded by the IARPA Metaphor Program)
3. Demonstration of language analysis and inference in automatic game player synthesis from natural language instructions
4. First prototype of analysis and inference framework for robot action synthesis with natural language instruction (joint work with Bielefeld University).

We continued our foundational research for the John Templeton Foundation, which in 2010 awarded us a competitive grant to explore the basic assumptions of artificial intelligence (AI) and cognitive science. Based on our long-standing effort in embodied semantics, we were able to formulate a sharp and precise definition of "actionability," the expected utility of an action, and produced arguments that this is much better than the traditional notion of "truth" as the core concept for cognitive science. This has led to the organization of an interdisciplinary Actionability Workshop, also sponsored by the Templeton Foundation, scheduled for April 2013.

April 20, 2012

AI Student Omoju Miller Attends TEDx Summit

Omoju Miller, a graduate student researcher in the AI Group, attended the inaugural TEDx Summit from April 16-20 in Doha, Qatar. The TEDx program sponsors local, independent events licensed by TED, the popular series of conferences on technology, entertainment, and design. The summit brings together 600 TEDx organizers from 90 countries to discuss ways to improve events. Last year, Miller organized TEDxEuclidAve, which focused on how to use technology and entrepreneurship to achieve social good.



AI

Publications

- M. BANSAL, J. DeNERO, AND D. LIN. Unsupervised Translation Sense Clustering. Proceedings of the North American Chapter of the Association for Computational Linguistics Human Language Technologies Conference (NAACL HLT 2010), Montreal, Canada, pp. 773-782, June 2012
- M. BANSAL AND D. KLEIN. Coreference Semantics from Web Features. Proceedings of the 50th Annual Meeting of the Association of Computational Linguistics (ACL 2012), Jeju, Korea, July 2012
- 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
- O. CULO AND G. DE MELO. Source-Path-Goal: Investigating the Cross-Linguistic Potential of Frame-Semantic Text Analysis. *it - Information Technology*, Vol. 54, Issue 3, pp. 147-152, May 2012
- O. CULO AND S. HANSEN-SCHIRRA. Treebanks in translation studies - the (CroCo) Dependency Treebank. Multilingual Corpora and Multilingual Corpus Analysis. Hamburg Studies in Multilingualism. In print, 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
- G. DE MELO, C. F. BAKER, N. IDE, R. J. PASSONNEAU, AND C. FELLBAUM. Empirical Comparisons of MASC Word Sense Annotations. Proceedings of the 8th Conference on International Language Resources and Evaluation (LREC 2012), Istanbul, Turkey, May 2012
- G. DE MELO, K. HOSE. Searching the Web of Data (Tutorial Abstract). Proceedings of the 35th European Conference on Information Retrieval (ECIR), Moscow, Russia, March 2012
- G. DE MELO AND G. WEIKUM. Constructing and Utilizing Wordnets using Statistical Methods. *Journal of Language Resources and Evaluation*, Vol. 46, Issue 2, pp. 287-311, June 2012
- G. DE MELO AND G. WEIKUM. UWN: A Large Multilingual Lexical Knowledge Base. Proceedings of the 50th Annual Meeting of the Association of Computational Linguistics (ACL 2012), Jeju, Korea, pp. 151-156, July 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
- C. J. FILLMORE. Encounters with Language. *Computational Linguistics*, Vol. 38, No. 4, pg.701-718, December 2012
- C. J. FILLMORE, R. R. LEE-GOLDMAN, AND R. RHOMIEUX. The FrameNet Construction. In *Sign-Based Construction Grammar*, H. C. Boas and I. A. Sag, eds., Center for the Study of Language and Information, 2012
- N. IDE, C. BAKER, C. FELLBAUM, AND R. PASSONNEAU. The MASC/MultiMASC Community Collaboration Project: Why You Should Be Involved and How. Proceedings of the Workshop on Collaborative Resource Development and Delivery at the 8th Conference on International Language Resources and Evaluation (LREC 2012), Istanbul, Turkey, May 2012
- P. KAY AND I. A. SAG. Cleaning up the Big Mess: Discontinuous Dependencies and Complex Determiners. In *Sign-Based Construction Grammar*, H. C. Boas and I. A. Sag, eds., Center for the Study of Language and Information, 2012
- P. KAY AND I. A. SAG. Discontinuous Dependencies and Complex Determiners. In *Sign-Based Construction Grammar*, H. C. Boas and I. A. Sag, eds., pp. 211-238, Center for the Study of Language and Information, 2012
- B. LOENNEKER-RODMAN AND S. NARAYANAN. Computational Models of Figurative Language. *Cambridge Encyclopedia of Psycholinguistics*, M. Spivey, M. Joannis, and K. McRae, eds., Cambridge University Press, 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
- 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
- R. PASSONNEAU, C. BAKER, C. FELLBAUM, AND N. IDE. The MASC Word Sense Corpus. Proceedings of the 8th Conference on International Language Resources and Evaluation (LREC 2012), Istanbul, Turkey, pp. 3025-3030, May 2012
- M. R. L. PETRUCK AND G. DE MELO. Precedes: A Semantic Relation in FrameNet. Proceedings of the Workshop on Language Resources for Public Security Applications at the 8th Conference on International Language Resources and Evaluation (LREC 2012), Istanbul, Turkey, pp.45-49, May 2012
- I. A. SAG, H. C. BOAS, AND P. KAY. Introducing Sign-Based Construction Grammar. In *Sign-Based Construction Grammar*, H. C. Boas and I. A. Sag, eds., pp. 1-28, Center for the Study of Language and Information, 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

TECHNICAL CONTENT

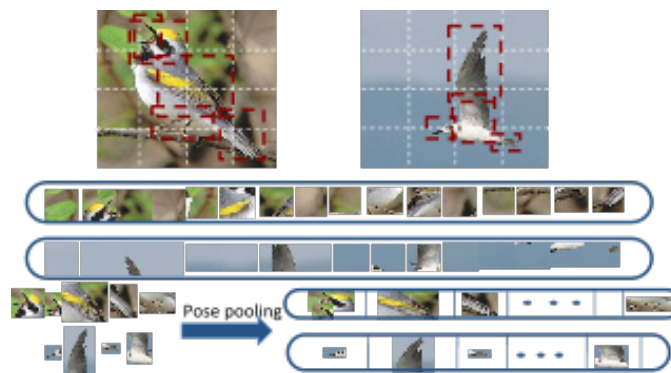
VISION

Fine-grained Recognition

Recognizing objects in fine-grained domains can be extremely challenging due to the subtle differences between subcategories. Discriminative markings are not only subtle but often highly localized, leading traditional object recognition approaches to struggle when dealing with the large pose variation often present in these domains.

The ability to normalize pose based on super-category landmarks can significantly improve models of individual categories when training data are limited. Previous methods have considered the use of volumetric or morphable models for faces and for certain classes of articulated objects. We consider methods that impose fewer representational assumptions on categories of interest, and exploit contemporary detection schemes which consider the ensemble of responses of detectors trained for specific pose-keypoint configurations. We develop representations for poselet-based pose normalization using both explicit warping and implicit pooling as mechanisms. Our method defines a pose normalized similarity or kernel function that is suitable for nearest-neighbor or kernel-based learning methods. This work has been reported at CVPR 2012; recently we have considered how to extend the method to rely on faster, more robust pose detectors.

Pose-normalization seeks to align training exemplars, either piecewise by part or globally for the whole object, effectively factoring out differences in pose and in viewing angle. Our effort factorizes the problem of pose-normalization into (i) the localization of semantic parts and (ii) learning an optimal description. For localization, we propose a part detector based on a strongly supervised variant of the state-of-the-art deformable part model. To describe the appearance of these parts or semantic "pooling regions," we utilize multiple-kernel-learning to select the best features for each subcategory. These methods are considered alongside previously proposed models for both the localization and description stages and rigorously studied in a comprehensive evaluation across multiple fine-grained datasets.




VISION

Representation Learning

Sequential Detection

Our group has focused on four core visual representation learning problems:

- Efficient learning of “deep” models
- Sharing sparse part descriptors
- Fine-scale segmentation models
- Time-optimized recognition



with deformable part models (Felzenszwalb et al., 2010a), showing significant speedup factors, but with a large decrease in task performance. We have described a new training framework that learns which sparselets to activate in order to optimize a discriminative objective, leading to larger speedup factors with no decrease in task performance. We first reformulate sparselets in a

general structured output prediction framework, then analyze when sparselets lead to computational efficiency gains, and lastly show experimental results on object detection and image classification tasks. Our experimental results demonstrate that discriminative activation substantially outperforms the previous reconstructive approach, which, together with our structured output prediction formulation, makes sparselets broadly applicable and significantly more effective.

Linear Support Vector Machines (SVMs) have become very popular in vision as part of state-of-the-art object recognition and other classification tasks but require high dimensional feature spaces for good performance. Deep learning methods can find more compact representations but current methods employ multilayer perceptrons that require solving a difficult, non-convex optimization problem. We propose a deep non-linear classifier whose layers are SVMs and which incorporates random projection as its core stacking element. Our method learns layers of linear SVMs recursively transforming the original data manifold through a random projection of the weak prediction computed from each layer. Our method scales as linear SVMs, does not rely on any kernel computations or nonconvex optimization, and exhibits better generalization ability than kernel-based SVMs. This is especially true when the number of training samples is smaller than the dimensionality of data, a common scenario in many real-world applications. The use of random projections is key to our method, as we show in the experiments section, in which we observe a consistent improvement over previous – often more complicated – methods on several vision and speech benchmarks.

Shared representations are highly appealing due to their potential for gains in computational and statistical efficiency. Compressing a shared representation leads to greater computational savings, but at the same time can severely decrease performance on a target task. Recently, sparselets (Song et al., 2012) were introduced as a new shared intermediate representation for multiclass object detection

Many object categories have rich fine-grained image structures, yet contemporary methods to discover categories from weakly labeled image sets have great difficulty segmenting such structures. While several methods offer precise segmentation of individual images, methods which learn an object model from several images have heretofore been insensitive to fine structures in an image, limiting the applicability of those methods. We have recently introduced a new method, which we call communal cuts, that provides a scalable, distributed, and integrable method for fine-grained co-segmentation of images. This method introduces higher-order coupling between pixel labels not only via joint foreground clustering, but also by explicitly coupling segmentation boundaries across images. This method both is theoretically well-motivated and empirically outperforms state-of-the-art methods. In addition, we introduce a new dataset for co-segmentation whose labelings are more detailed than those of existing ones, enabling benchmarking and learning for high-quality segmentation methods. This is work in review by Evan Shelhamer, Stefanie Jegelka, and Trevor Darrell.



TECHNICAL CONTENT

In a large visual multi-class detection framework, the timeliness of results can be crucial. Our method for timely multi-class detection aims to give the best possible performance at any single point after a start time; it is terminated at a deadline time. Toward this goal, we formulate a dynamic, closed-loop policy that infers the contents of the image in order to decide which detector to deploy next. In contrast to previous work, our method significantly diverges from the predominant greedy strategies and is able to learn to take actions with deferred values. We evaluate our method with a novel timeliness measure,

computed as the area under an Average Precision vs. Time curve. Experiments are conducted on the PASCAL VOC object detection dataset. If execution is stopped when only half the detectors have been run, our method obtains 66 percent better AP than a random ordering, and 14 percent better performance than an intelligent baseline. On the timeliness measure, our method obtains at least 11 percent better performance. Our method is easily extensible, as it treats detectors and classifiers as black boxes and learns from execution traces using reinforcement learning.

Perceptual Security

Together with collaborators from Professor Dawn Song's group at UC Berkeley, we have investigated new models for cybersecurity that draw on models for human and machine perception. We revisit user interface (UI) security attacks, such as clickjacking, from a perceptual perspective and develop an innovative defense as well as identify novel attacks. We develop a computer vision-based defense for defending against common clickjacking attacks on a perceptual level. Our system is open source and ensures contextual integrity by analyzing OS screenshots, cursor movements, and clicking actions. Our perceptual view on UI security attacks also helps identify new attacks on UI security. We develop five attacks that bypass current defenses. Our attacks are powerful with a 100 percent success rate in some cases. However, they only scratch the surface of possible perceptual attacks on UI integrity, and we posit that a number of attacks are possible with a comprehensive study of human perception. Finally, we argue that, due to the complex nature of human perception, defending against such attacks is challenging and requires further research taking user perception and new computer vision techniques into account.

This work is under review and was led by Erik Rodner of the Vision Group.

In further work, we have investigated visual phishing defenses. Phishing attacks are an increasingly pervasive threat to the privacy and security of Web users. Such attacks commonly rely on fooling victims into visually perceiving an attacker-controlled Web page as content from a legitimate source. We demonstrate that existing automated phishing detection systems will fail to detect phishing pages that have high visual similarity to the source pages, and examine the use of computer vision techniques in detecting such visual perception-based attacks. In particular, we show how computer vision can be used to detect exact copies or modified versions of logos and other trademarks of popular, high-value phishing targets. Some of our vision-based phishing detection methods are suitable for use in a local browser extension, in addition to the possibility of enhancing existing cloud blacklists. This work is in preparation, with Erik Rodner and Jeff Donahue from the ICSI vision group.

VISION

Robotic vision

To perform useful tasks in everyday human environments, robots must be able to both understand and communicate the sensations they experience during haptic interactions with objects. Toward this goal, we augmented the Willow Garage PR2 robot with a pair of SynTouch BioTac sensors to capture rich tactile signals during the execution of four exploratory procedures on 60 household objects. In a parallel experiment, human subjects blindly touched the same objects and selected binary haptic adjectives from a predetermined set of 25 labels. We developed several machine-learning algorithms to discover the meaning of each adjective from the robot's sensory data. The most successful algorithms were those that intelligently combine static and dynamic components of the data recorded during all four exploratory procedures. The best of our approaches produced an average adjective classification F1 score of 0.77, a score higher than that of an average human subject. We are presently extending this work to perform visual assessment of haptic properties, using a joint visuo-haptic database.

In work in review, we develop a grasp affordance approach that combines image-based category level detection methods with 3D point cloud data. Purely 3D data-based grasp methods have been widely used for robotics tasks but most 3D grasp approaches lack the ability to grasp an object at a certain part, e.g., a cup on a handle only and not from inside. They are also hard to apply on flat objects on a table plane. Image-only-based methods, e.g., the Deformable Parts Model, showed promising results for object classification and pose estimation but are often too inaccurate to be applied on physical real world objects. The advantage of a combination of both 3D data and a 2D grasp point estimation pipeline enables the robot to grasp flat objects, to grasp certain parts of a 3D point cloud, and to generalize over object instances. We performed experiments using a set of household objects and will compare our combined method to a 2D and 3D base line approach. As a testbed within

our experiments served a PR2 robot from Willow Garage, equipped with arms, grippers, as well as 2D and 3D sensors. This is ongoing work with Daniel Göhring, Lorenzo Riano, Hyun Oh Song, Trevor Darrell.

Further, we have investigated grounding of spatial relations for Human-Robot Interaction. Natural language understanding is a key requirement to have humans and robots naturally interact with each other. However, in order to understand language, a robot needs to create the appropriate grounding between symbols in a sentence and the physical world as perceived by its sensors. Based on previous works that highlight the relevance of spatial relations in human-robot interactions and in the idea that models learnt automatically can lead to better performance, greater flexibility and adaptation capability, we have developed a system that learns models for spatial prepositions and object recognition in order to understand statements that refer to objects (nouns), their spatial relationship (prepositions), and to execute different commands (verbs) upon request. The proposed method is inherently compositional, and allows the robot to understand statements such as "Move the cup close to the robot to the area in front of the plate and behind the tea_box". To test the feasibility of our proposed method, it has been deployed and tested on a PR2 robot. In this setting the robot is capable of interpreting sentences that include references to objects, solving complex spatial queries and showing the results of grounding by executing the appropriate action. This is work in review, with Sergio Guadarrama, Lorenzo Riano, Dave Golland, Daniel Gohring, Yangqing Jia, Dan Klein, Pieter Abbeel and Trevor Darrell.



TECHNICAL CONTENT

VISION

Domain Adaptation

Most successful object classification and detection methods rely on classifiers trained on large labeled datasets. However, for domains where labels are limited, simply borrowing labeled data from existing datasets can hurt performance, a phenomenon known as "dataset bias." We propose a general framework for adapting classifiers from "borrowed" data to the target domain using a combination of available labeled and unlabeled examples. Specifically, we show that imposing smoothness constraints on the classifier scores over the unlabeled data can lead to improved adaptation results. Such constraints are often available in the form of instance correspondences, e.g., when the same object or individual is observed simultaneously from multiple views, or tracked between video frames. In these cases, the object labels are unknown but can be constrained to be the same or similar. We propose techniques that build on existing domain adaptation methods by explicitly modeling these relationships, and demonstrate empirically that they improve recognition accuracy in two scenarios, multi-category image classification and object detection in video.

Object recognition has made significant advances in recent years. However, algorithm over-fitting causes degraded performance when there is a mismatch between training and test data. Domain adaptation offers both supervised and unsupervised solutions to the data mismatch problem by adapting a trained

classifier to the target (test) data. The classic domain adaptation paradigm considers the world of images to be separated into discrete domains where adaptation can be learned between any pair of domains. This definition of domain adaptation is particularly applicable to the discretized settings of natural language processing and speech recognition. In vision, however, our input space cannot be naturally separated into discrete domains. A more realistic view of the visual world is to realize that visual images are drawn from a continuous manifold of domains. Recent work has introduced the idea of placing discrete domains onto a manifold. In this work we extend the notion of domains as points on a manifold to consider each image as potentially being drawn from a different domain and hence a different point on the manifold. We use a rank constraint on the manifold subspaces to restrict the possible space and use class and instance constraints to provide smoothness (locality) constraints on the subspace positions. Using our algorithm we are able to learn a domain manifold that provides a domain invariant representation which is used for classification of a test point. This is work in progress, by Judy Hoffman, Erik Rodner, Jeff Donahue, Kate Saenko, and Trevor Darrell.



Selected References

V. CHU, I. McMAHON, L. RIANO, C. G. McDONALD, Q. HE, J. MARTINEZ PEREZ-TEJADA, M. ARRIGO, N. FITTER, J. C. NAPPO, T. DARRELL, AND K. J. KUCHENBECKER. Using Robotic Exploratory Procedures to Learn the Meaning of Haptic Adjectives. Proceedings of the 2013 IEEE International Conference on Robotics and Automation, Karlsruhe, Germany, May 20, 2013.

J. DONAHUE, J. HOFFMAN, E. RODNER, K. SAENKO, AND T. DARRELL. Semi-Supervised Domain Adaptation with Instance Constraints. Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR 2013), Portland, Oregon, June 2013.

R. GIRSHICK, H. O. SONG, AND T. DARRELL. Discriminatively Activated Sparselets. Proceedings of the International Conference on Machine Learning (ICML 2013), Atlanta, Georgia, June 2013.

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.

I. McMAHON, V. CHU, L. RIANO, C. G. McDONALD, Q. HE, J. M. PEREZ-TEJADA, M. ARRIGO, N. FITTER, J. C. NAPPO, T. DARRELL, AND K. J. KUCHENBECKER. Robotic Learning of Haptic Adjectives Through Physical Interaction. Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems, Vilamoura, Portugal, October 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.

N. ZHANG, R. FARRELL, AND T. DARRELL. Pose Pooling Kernels for Sub-Category Recognition. Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR 2012), Providence, Rhode Island, pp. 3665-3672, June 2012.



VISION Publications

- T. ALTHOFF, H. O. SONG, AND T. DARRELL. Detection Bank: An Object Detection Based Video Representation for Multimedia Event Recognition. Proceedings of the ACM Multimedia Conference (ACMMM '12), Nara, Japan, October 2012.
- R. BOEHME AND M. KIRCHNER. Counter-Forensics: Attacking Image Forensics. In Digital Image Forensics, H. T. Sencar and N. D. Memon, eds., pp. 327-366, Springer, 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, November 2012. Recipient of a best paper honorable mention award.
- 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.
- 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.
- 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.
- Y. JIA, C. HUANG, AND T. DARRELL. Beyond Spatial Pyramids: Receptive Field Learning for Pooled Image Features. Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR 2012), Providence, Rhode Island, pp. 3370-3377, June 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 AND D. GHOSAL (EDS.). Information Hiding. 14th International Conference, Information Hiding 2012, Berkeley, California, Revised Selected Papers, Lecture Notes in Computer Science Vol. 7692, May 2012.
- S. MILLER, J. VAN DEN BERG, M. FRITZ, T. DARRELL, K. GOLDBERG, AND P. ABBEEL. A Geometric Approach to Robotic Laundry Folding. The International Journal of Robotics Research, Vol. 31, No. 2, pp. 249-267, February 2012.
- S. PFENNIG AND M. KIRCHNER. Spectral Methods to Determine the Exact Scaling Factor of Resample Digital Images. Invited paper, "Image Forensics Techniques" special session, proceedings of the 5th International Symposium on Communications, Control, and Signal Processing (ISCCSP), Rome, Italy, May 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.
- 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.
- S. VIRTANEN, Y. JIA, A. KLAMI, AND T. DARRELL. Factorized Multi-Modal Topic Model. Proceedings of the Conference on Uncertainty in Artificial Intelligence (UAI), Catalina Island, California, August 2012.
- Y. XIONG, K. SAENKO, T. DARRELL, AND T. ZICKLER. From Pixels to Physics: Probabilistic Color De-rendering. Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), Providence, Rhode Island, June 2012.
- N. ZHANG, R. FARRELL, AND T. DARRELL. Pose Pooling Kernels for Sub-Category Recognition. Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR 2012), Providence, Rhode Island, pp. 3665-3672, June 2012.
- W. ZHANG, S. X. YU, AND S.-H. TENG. Power SVM: Generalization with Exemplar Classification Uncertainty. Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR 2012), Providence, Rhode Island, pp. 2144-2151, June 2012.

TECHNICAL CONTENT

ALGORITHMS

Population Genomics

Genetic studies of diseases involve the collection and sequencing of many individuals. The analysis of the sequence data involves several computational challenges, including the identification of genetic variations and the assignment of those to specific traits. We developed a set of computational methods for the identification of specific genetic variants, called Copy Number Variants (CNVs), where a stretch of the genome in certain individuals may be deleted or copied into different places of the genome. Such variants are extremely hard to detect using current computational methods. However, there is evidence that these variants are occasionally related to diseases such as autism, cancer, and others.

The analysis of multiple individuals simultaneously requires the development of methods and resources that will allow for the genomic material of one individual to be related to another, and for those to be related to world's population. We developed a set of methods that facilitate such analysis, including methods that can identify the geographic origins of an individual solely based on his or her genetic data. In one of the projects, published in the journal *Nature Genetics*, we show that the origin of an individual with European ancestry can be detected with striking accuracy: a median error of 250 kilometers (see Figure 1). Our method is the first to suggest a random generative model for spatial genetic data. Such a model describes the "typical" genome at each position (x,y) in the map.

We also contributed to the analysis of the international 1000 Genomes project (*Nature*, 2012), which is an international effort of a large number of groups to provide the entire genome

sequences of more than a thousand people from different world populations. Researchers across the world later use these sequences for the analysis of disease-related data that they generate in their lab.




Figure 1 – The geographic mapping of genetic data of individuals of European ancestry. The input to the algorithm is the genome sequence of each individual, and the output is an (x,y) position. The colors in the figure correspond to the different countries of origin, as well as the naming (e.g., ES = Spain, IT = Italy, FR = France).

ALGORITHMS

Brain Networks

With the currently available data and that expected in the near future from the Human Connectome Project, the development of tools and skills for its analysis will have significant impact in both clinical and research settings. The primary aim of this project is to develop tools for the study of Brain Networks. In collaboration with UCSF experts in the area of brain imaging and Brain Network reconstruction (Roland Henry, Pratik Mukherjee and Srikanth Nagarajan in the UCB/UCSG joint Bioengineering group and Karl Young and Norbert Shuff at the San Francisco Va Medical Center) we study a variety of brain networks arising from multiple imaging modalities (MEG, fMRI and DTI/HARDI) and develop tools to both understand pathologies and make predictions. Current projects focus on Agenesis of the Corpus Callosum (a birth defect), Multiple Sclerosis, Alzheimer's disease and Traumatic Brain Injury.

This research is focused on developing robust localized network measures. This project is based on the hypothesis that one can construct robust localized network measures that will provide valuable diagnostic tools for neural disorders of the brain. This is needed because many disorders and the specific

conditions caused by those disorders are likely to depend on network connectivity issues in specific regions of the brain, which can be missed by global measures, the focus of most studies on Brain Networks to date. These measures include localized versions of many well-known global measures such as assortativity, efficiency, degree distributions, modularity and small-worldness, obtained by restricting to localized subnetworks of the brain network. It is analyzing localized versions of other less well known robust measures, such as flow based measures and expansion coefficients. Also, it is developing localized versions of robust Bayesian methods by modifying priors to include positional information.

To develop these measures and create a baseline for future work it is rigorously testing these methods on a suite of functional and structural Brain Networks measured on young adults, using a variety of statistical techniques. In addition, it is developing fast and accurate algorithms for creating reference networks of the bootstrap and Bayesian methods. The outcome of this project will be a publicly available suite of robust localized network analysis tools that will be available to researchers and clinicians.

Manipulation in Data Centers

Cloud computing is becoming a major paradigm for modern computation and information technology. One of the key issues in developing cloud computing systems is the development of fast, robust and efficient resource allocation mechanisms for sharing the resources in the cloud among a large number of heterogeneous users under a variety of constraints. One emerging approach is the use of price mechanisms. Price mechanisms are the foundation of much of the world's economy and are remarkably robust and efficient for allocating resources, but they have several potential weaknesses in the cloud computing setting. These include the speed and robustness of convergence to equilibrium and the manipulability of the equilibrium. These issues are less significant in slow moving homogeneous markets,

which are common in the "economic" world, but pose a problem in the fast changing extremely heterogeneous markets that underlie clouds of computing resources.

This project takes a principled approach to understanding and mitigating these problems with price mechanisms for cloud computing. This project uses analytic analyses combined with simulations to provide a detailed understanding of price mechanisms for cloud computing and aid in the development of efficient price mechanisms which converge rapidly and robustly and both prevent and mitigate the losses due to manipulations.



TECHNICAL CONTENT

ALGORITHMS

Studies of Human Disease

We are not merely developing theoretical methods – we also apply these methods to the analysis of specific diseases. We are collaborating with international groups that study specific diseases, including atherosclerosis, coronary artery disease, leukemia, and non-Hodgkin Lymphoma. In the past year we were able to detect multiple genetic variants that are associated with leukemia and with lipid concentrations. It is well known that lipid concentrations are correlated with cardiovascular diseases such as atherosclerosis, and therefore such discoveries may impact medicine and particularly personalized medicine.

A few of the disease-specific studies are ongoing multiple-year projects. In one study, we are working on the first whole-genome sequencing study of Non-Hodgkin lymphoma, in which we have sequenced 250 cases (individuals carrying Follicular lymphoma) in order to compare their genomes to the 1000 genomes sequences described above. Our initial results are promising: we found 31 regions where rare genetic variants are more abundant in Follicular Lymphoma patients compared to the general population. These findings are being validated in the lab of our collaborators.

Metagenomics

The number of bacteria and viruses living within our body is an order of magnitude larger than our own cells. Therefore, in order to study our own genetics and relate the genetic variants to diseases, it is important to characterize the genetic diversity of the microbiome living with us and the variation of bacterial species as a function of disease and other traits. For instance, it has been shown that the composition of gut bacteria is different between individuals with Crohn's disease and healthy individuals. In order to study these bacteria, biologists sample a part of the body and sequence the entire bacterial community.

It is then a challenging computational task to separate the sequence data into species.

We have developed methods that try to answer this challenge, both for the analysis of HIV, i.e., assessing the types of HIV developing under drug resistance, and for the more general question of de-convoluting the bacterial sequences into separate bacteria. We were the first to propose a method that provides an analysis of multiple microbial samples and correlates bacterial sequencing features with disease, gender, age, and country of origin.

Group Highlights

In recent years, sequencing and genotyping capacity has increased dramatically, leading to exciting discoveries in human genetics regarding the relations between genes and diseases. The extremely large amounts of data generated by these technologies set computational and statistical challenges; a large part of the Algorithms Group's work involves searching for methods that cope with these challenges.

The group has been mainly focusing on the development of computational methods for the analysis of high-throughput genomics data, particularly in relation to disease and to population studies. These studies pose many computational problems stemming from the fact that the sequencing and genotyping technologies provide noisy and incomplete data, and since in some cases the information about the genotyped subjects is limited.

ALGORITHMS

Publications

Y. BARAN AND E. HALPERIN. Joint Analysis of Multiple Metagenomic Samples. *PLOS Computational Biology*, accepted for publication, 2012

Y. BARAN, B. PASANIUC, S. SANKARAMAN, D. G. TORGERSON, C. GIGNOUX, C. ENG, W. RODRIGUEZ-CINTRON, R. CHAPELA, J. G. FORD, P. C. AVILA, J. RODRIGUEZ-SANTANA, E. G. BURCHARD, AND E. HALPERIN. Fast and Accurate Inference of Local Ancestry in Latino Populations. *Bioinformatics*, Vol. 28, Issue 10, pp. 1359-1367, May 2012

A. EFROS AND E. HALPERIN. Haplotype Reconstruction Using Perfect Phylogeny and Sequence Data. *Proceedings of the Second Annual Satellite Workshop on Massively Parallel Sequencing at the 16th Annual International Conference on Research in Computational Molecular Biology (RECOMB 2012)*, Barcelona, Spain. Published in *BMC Bioinformatics*, Vol. 13, Suppl. 6:S2, April 2012

R. KARP. Theory of Computation as an Enabling Tool for the Sciences. Presented at the Turing Centennial Celebration, Princeton, New Jersey. Also presented at the 9th Annual Conference on Theory and Applications of Models of Computation (TAMC 2012), Turing Centenary Meeting, Beijing, China, May 2012

B. KIRKPATRICK, Y. RESHEF, H. FINUCANE, H. JIANG, B. ZHU, AND R. M. KARP. Comparing Pedigree Graphs. *Journal of Computational Biology*, 2012

J. NISHIMURA AND E. FRIEDMAN. Probabilistic Convergence Guarantees for Type-II Pulse-Coupled Oscillators. *Physical Review E*, Vol. 86, Issue 2, August 2012

R. ROZOV, E. HALPERIN, AND R. SHAMIR. MGMR: Leveraging RNA-Seq Population Data to Optimize Expression Estimation. *Proceedings of the Second Annual Satellite Workshop on Massively Parallel Sequencing at the 16th Annual International Conference on Research in Computational Molecular Biology (RECOMB 2012)*, Barcelona, Spain. Published in *BMC Bioinformatics*, Vol. 13, Suppl. 6:S2, April 2012

R. SHARAN AND R. M. KARP. Reconstructing Boolean Models of Signaling. *Proceedings of the 16th Annual International Conference on Research in Computational Molecular Biology (RECOMB 2012)*, Barcelona, Spain, pp. 261-271, April 2012

S. L. SLAGER, C. F. SKIBOLA, ..., E. HALPERIN, ET AL. Common Variation at 6p21.31 (BAK1) Influences the Risk of Chronic Lymphocytic Leukemia. *Blood*, Vol. 120, No. 4, pp. 843-846, July 2012

Z. WANG, F. HORMOZDIARI, W.-Y. YANG, E. HALPERIN, AND E. ESKIN. CNVeM: Copy Number Variation Detection Using Uncertainty of Read Mapping. *Proceedings of the 16th Annual International Conference on Research in Computational Molecular Biology (RECOMB 2012)*, Barcelona, Spain, pp. 326-340, April 2012

W.-Y. YANG, J. NOVEMBRE, E. ESKIN, AND E. HALPERIN. A Model-Based Approach for Analysis of Spatial Structure in Genetic Data. *Nature Genetics*, Vol. 44, pp. 725-731, May 2012. Also Presented at the Annual Meeting of the American Society of Human Genetics, San Francisco, California, November 2012

O. ZAGORDI, A. TOPFER, S. PRABHAKARAN, V. ROTH, E. HALPERIN, AND N. BEERENWINKEL. Probabilistic Inference of Viral Quasispecies Subject to Recombination. *Proceedings of the 16th Annual International Conference on Research in Computational Molecular Biology (RECOMB 2012)*, Barcelona, Spain, pp.342-354, April 2012

Selected references

Y. BARAN, B. PASANIUC, S. SANKARAMAN, D. G. TORGERSON, C. GIGNOUX, C. ENG, W. RODRIGUEZ-CINTRON, R. CHAPELA, J. G. FORD, P. C. AVILA, J. RODRIGUEZ-SANTANA, E. G. BURCHARD, AND E. HALPERIN. Fast and Accurate Inference of Local Ancestry in Latino Populations. *Bioinformatics*, Vol. 28, Issue 10, pp. 1359-1367, May 2012.

A. DEMIRKAN, ET AL. Genome-wide association study identifies novel loci associated with circulating phospho- and sphingolipid concentrations. *PLoS Genetics*, Vol. 8, No. 2: e1002490, February 2012.

A. EFROS AND E. HALPERIN. Haplotype Reconstruction Using Perfect Phylogeny and Sequence Data. *Proceedings of the Second Annual Satellite Workshop on Massively Parallel Sequencing at the 16th Annual International Conference on Research in Computational Molecular Biology (RECOMB 2012)*, Barcelona, Spain. Published in *BMC Bioinformatics*, Vol. 13, Suppl. 6:S2, April 2012.

R. ROZOV, E. HALPERIN, AND R. SHAMIR. MGMR: Leveraging RNA-Seq Population Data to Optimize Expression Estimation. *Proceedings of the Second Annual Satellite Workshop on Massively Parallel Sequencing at the 16th Annual International Conference on Research in Computational Molecular Biology (RECOMB 2012)*, Barcelona, Spain. Published in *BMC Bioinformatics*, Vol. 13, Suppl. 6:S2, April 2012.

S.L. SLAGER, C.F. SKIBOLA, ..., E. HALPERIN, ET AL. Common variation at 6p21.31 (BAK1) influences the risk of chronic lymphocytic leukemia. *Blood*, Vol. 120, No. 4, pp 843-846, July 2012.

Z. WANG, F. HORMOZDIARI, W.-Y. YANG, E. HALPERIN, AND E. ESKIN. CNVeM: Copy Number Variation Detection Using Uncertainty of Read Mapping. *Proceedings of the 16th Annual International Conference on Research in Computational Molecular Biology (RECOMB 2012)*, Barcelona, Spain, pp. 326-340, April 2012.

W.-Y. YANG, J. NOVEMBRE, E. ESKIN, AND E. HALPERIN. A model-based approach for analysis of spatial structure in genetic data. *Nature Genetics*, Vol. 44, pp. 725-731, May 2012.

O. Zagordi, A. Topfer, S. Prabhakaran, V. Roth, E. Halperin, and N. Beerenwinkel. Probabilistic Inference of Viral Quasispecies Subject to Recombination. *Proceedings of the 16th Annual International Conference on Research in Computational Molecular Biology (RECOMB 2012)*, Barcelona, Spain, pp.342-354, April 2012.

1000 GENOMES PROJECT CONSORTIUM. An integrated map of genetic variation from 1,092 human genomes. *Nature*, Vol 491, pp. 56-65, November 2012.



TECHNICAL CONTENT

ARCHITECTURE

Group Highlights

Our collaboration with MIT and Micron on photonically connected DRAM chips has resulted in several test chip fabrications, including the first monolithically integrated transmitter developed by our colleagues at MIT and presented at ISSCC 2013.

We have also successfully fabricated the first monolithically integrated photonics microprocessor incorporating an integrated photonic link to memory in IBM 45nm technology, and subsequently designed and taped out two further photonics microprocessors.

In our resilient vector-thread architecture work, we have taped out a further design in 28m technology.

Vector-Thread Architectures

In earlier work at MIT, Professor Krste Asanovic's team developed the Scale vector-thread architecture and processor prototype, which combines data-level and thread-level parallel execution models in a single unified architecture. Maven is the second-generation vector-thread architecture, designed to scale up to hundreds of execution "lanes", and with the goal of providing very high throughput at low energy for a wide variety of parallel applications. Maven is based on a new compact lane design, which is replicated to yield a "sea-of-lanes" execution substrate. At run-time, lanes are ganged together to form variable-sized vector-thread engines, sized to match application needs. Completing our work on Maven, we were invited to expand on our ISCA 2011 paper exploring the efficiency versus programmer productivity tradeoffs of various data-parallel accelerators for an ACM Transactions on Computer Systems paper to appear in 2013.

We have now started a new data-parallel architecture project, Hwacha, which we believe will provide a further improvement over the earlier vector-thread architectures by moving more complexity into software. In collaboration with students at the Parallel Computing Laboratory, or Par Lab, in the Computer Science Division at UC Berkeley, we have demonstrated a working auto-vectorizing Python compiler that can generate high-performance vector machine code from Python source code. We are working to improve the microarchitecture and the vector compiler, which is based on LLVM. Our hope is to provide even greater energy efficiency while retaining the flexibility of our earlier designs.

In 2011, we had taped out an early Hwacha vector-thread prototype microprocessor, Raven-1, in a pre-production 28nm process from ST microelectronics, including features to provide fault resiliency when operating at very low voltages. The fabrication encountered several setbacks, delaying the arrival of functioning parts and causing several problems with the finished parts. In recent testing, we have been able to demonstrate instructions executing on the processor core and are gathering data on SRAM and processor resiliency and energy efficiency.

We have subsequently taped out a second variant of this resilient vector architecture, Raven-2, in a near-production version of the 28nm process. We have now received parts and are in the process of packaging the chip for testing.

We are currently working toward a quad-core version, Raven-3, for tapeout in the summer, which will incorporate multiple on-chip DC-DC convertors that use a novel technique to skip the last stage of conventional DC regulation and instead allow the processor to operate with a rippling supply voltage to boost energy efficiency. Raven-3 will represent the most ambitious design undertaken by our group to date, as it will include four vector cores operating at independent supply voltages and clock frequencies connected via a coherent last-level cache.

ARCHITECTURE

Monolithic Silicon Photonics


In collaboration with the MIT Center for Integrated Photonic Systems, architecture researchers are exploring the use of silicon photonics to meet the bandwidth needs of future manycore processors. Projected advances in electrical signaling seem unlikely to fulfill memory bandwidth demands at feasible pinouts and power consumptions. Monolithic silicon photonics, which integrates optical components with electrical transistors in a conventional CMOS process, is a promising new technology that could provide both large improvements in achievable interconnect bandwidth and large reductions in power, but with much lower cost than alternative approaches that require separate fabrication of photonic components.

In earlier work, we explored how to use photonics to interconnect many processors to memory controllers using an opto-electrical global crossbar. Our results showed that aggregate throughput can be improved by about 8-10 compared to an optimized purely electrical network. We have also explored the use of photonics to connect DRAM chips to processor chips directly. Our initial designs promise large improvements in bandwidth and energy efficiency, provided the required photonic devices can be fabricated successfully in a DRAM process. We received a three-year award from DARPA to fund a collaboration with Micron Technology, which has been fabricating prototypes of our photonic devices in their DRAM technology and working with us to explore alternative approaches for connecting DRAMs to microprocessors photonically.

A major milestone was achieved by the group this year when our MIT colleagues presented the first fully functional monolithically integrated photonic transmitter in a paper and demonstration session at the ISSCC conference in February 2013.

Another milestone was our fabrication of the first monolithically integrated photonics microprocessor (EOS14), incorporating one of our vector processors with a photonic memory system interconnect in a standard IBM 45nm SOI technology. When running in an initial constrained test package, the processor in EOS14 successfully ran tests at over 900MHz, where the clock frequency was limited by our test infrastructure. The chip is expected to run considerably faster once re-packaged. A challenge in our experimental photonics work is the immaturity of commercial packaging for high-density photonics interconnect.

EOS14 contained a single vector core, but we have since taped out two further photonic microprocessors (EOS16 and EOS18), both containing dual vector cores with coherent caches as well as improved versions of the photonics interconnect. EOS16 has returned from manufacturing and is awaiting packaging. EOS18 contains several improvements to processor frequency and energy consumption and is expected to exceed 1.5GHz under typical conditions. Our MIT and Micron partners have also successfully fabricated a number of test chips in Micron's DRAM process.





TECHNICAL CONTENT

ARCHITECTURE

Other Collaborations

The Architecture Group works closely with Par Lab, using the software tools and parallel applications developed there to evaluate new architectural ideas.

As part of our VLSI effort, we have been collaborating on a Department of Education-funded project, "Isis" (Infrastructure for Synthesis with Integrated Simulation), at the Berkeley Wireless Research Center, to further develop our VLSI tools to provide a parameterized generation of optimized custom many-core processors. In Isis, we developed a new hardware design language called "Chisel" (Constructing Hardware in a Scala Embedded Language), which extends the powerful general-purpose host language Scala with hardware design libraries. A paper on Chisel appeared in the 2012 Design Automation Conference.

Chisel was first used in teaching the graduate VLSI class in the fall 2011 semester, and has subsequently been used in multiple undergraduate and graduate classes and now forms the primary hardware description language for all of our hardware design work at Berkeley. We have made Chisel available open-source under a BSD license, and we ran a "bootcamp" in summer 2012 for interested outsiders, with around 50 academic and industrial researchers attending. We have subsequently hired a staff engineer to help support Chisel, and have been working to improve the language, standard libraries, and examples to help with Chisel uptake in the broader VLSI community.

In 2011, we developed a new open-source instruction set architecture, called RISC-V, which has formed the basis of all of our research processor work and which is also used for education in multiple classes at UC Berkeley. Last year, we publicly released a variety of educational RISC-V processors written in Chisel in the "Sodor" repository, for use in undergraduate computer architecture classes. We have now begun reaching to other major universities that have expressed interest in using our RISC-V infrastructure in education and research. Recent highlights including booting an initial Linux port, undertaken by two enterprising undergraduate researchers, on RISC-V.

A major development on campus has been the winding down of the UC Berkeley Par Lab project, which officially ends in May 2013, and the start of the five-year ASPIRE project, led by Professor Krste Asanovi and initially funded by the DARPA PERFECT program together with several industrial affiliates. ASPIRE recognizes the shift from transistor-scaling-driven performance improvements to a new post-scaling world, where whole-stack co-design is the key to improved efficiency. Building on the success of Par Lab, ASPIRE aims to use deep hardware and software co-tuning to achieve the highest possible performance and energy efficiency for future mobile and rack computing systems. ASPIRE will extend our work in vector-thread architecture to more specialized accelerators, and will also support continued development of RISC-V-based processor and the Chisel infrastructure.

References

J. BACHRACH, H. VO, B. RICHARDS, Y. LEE, A. WATERMAN, R. AVIZIENIS, J. WAWRZYNEK, AND K. ASANOVIC. Chisel: Constructing Hardware in a Scala Embedded Language. Proceedings of the Design Automation Conference (DAC-2012), San Francisco, California, June 2012.

C. BATTEN, A. JOSHI, V. STOJANOVIC, AND K. ASANOVIC. Designing Chip-Level Nanophotonic Interconnection Networks. IEEE Journal on Emerging and Selected Topics in Circuits and Systems, Vol. 2, No. 2, pp. 137-153, June 2012.


S. BEAMER, A. BULUÇ, K. ASANOVIC, AND D. PATTERSON. Distributed Memory Breadth-First Search Revisited: Enabling Bottom-Up Search. Proceedings of the Workshop on Multithreaded Architectures and Applications (MTAAP-2013), at the International Parallel & Distributed Processing Symposium (IPDPS-2013), Boston, Massachusetts, May 2013.

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.

Y. LEE, R. KRASHINSKY, V. GROVER, S. W. KECKLER, AND K. ASANOVIC. Convergence and Scalarization for Data-Parallel Architectures. Proceedings of the International Symposium on Code Generation and Optimization (CGO 2013), Shenzhen, China, February 2013.

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. ACM Transactions on Computer Systems, 2013.

NEWS



Roberto Pieraccini Appointed New Director of ICSI

ICSI has appointed Roberto Pieraccini as its new director. A renowned leader in science and technology, Pieraccini brings to ICSI expertise in both academic and corporate research as well as enterprise business leadership.

"The International Computer Science Institute brings together incredibly talented world-class computer scientists, including some of the most widely cited senior and accomplished younger researchers, to tackle some of our society's most pressing intellectual challenges," said Pieraccini. "I am honored to continue in the footsteps of my predecessor Nelson Morgan, who fostered a stimulating international environment and helped ICSI grow as one of the preeminent centers of theoretical and applied research in computer science."

Algorithms Leader Richard Karp to Be Director of Simons Institute

Professor Richard Karp, leader of the ICSI Algorithms Group, will be the founding director of a theoretical computer science institute to be established at UC Berkeley with a \$60 million gift from the Simons Foundation. The Simons Institute for the Theory of Computing will bring together top computer theorists and researchers from around the world to address challenges in areas that include social sciences, biology, physics, and economics. ICSI is among several research centers to be affiliated with the institute, which will begin operations in July.



ICSI Appoints New Members to Board of Trustees

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 Friday, 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 Leader Scott Shenker Appointed Chief Scientist of ICSI

ICSI has appointed Scott Shenker as its first chief scientist. Shenker, who has led ICSI's Networking Group since 1998, will help set its research agenda and represent the interests of its researchers. He is the first to serve as chief scientist of ICSI.

"We established the role of chief scientist at ICSI to ensure the Institute continues with our research excellence and expands our worldwide scientific relevance in today's quickly evolving technological landscape," said Director Roberto Pieraccini. "And Scott Shenker, a true visionary, a role model in the computer science community and one of the pillars of ICSI, is ideally qualified to take on this role."

Bro 2.1 Released

Bro 2.1 was publicly released August 29, shortly after the 2012 Bro Exchange, a meeting that brought together Bro users to exchange thoughts on and experiences with deploying the system. Bro is an open-source network security monitor developed by a team of researchers and engineers at ICSI




and the National Center for Supercomputing Applications at the University of Illinois. Bro monitors network traffic at major universities, large research labs, supercomputing centers, and open-science communities around the country. It has been the cornerstone of network security at the Lawrence Berkeley National Laboratory since the 1990s. Recent installations include an OpenFlow-based Bro cluster setup at Indiana University. The University of Utah's Center for High Performance Computing is also currently investigating how to integrate the monitor into its new security architecture. The new 2.1 release introduces extensive support for IPv6, tunnel decapsulation, a new input framework for integrating external information in real-time into the processing, and two new experimental log output formats.

Chairman of ICSI's Board Accepts Position at Google

Prabhakar Raghavan, chairman of ICSI's Board of Trustees, has accepted a position as the vice president of strategic technologies at Google. Raghavan was formerly the head of Yahoo! Labs, where he worked for the last seven years.

MIT Releases Book by Director Roberto Pieraccini

The Voice in the Machine, the latest book by ICSI Director Roberto Pieraccini, is now available from MIT Press. The book examines the evolution of human-computer speech interaction capabilities over the past six decades.



"The Voice in the Machine is a fascinating look at how speech recognition technology has developed from its early beginnings to today's widespread deployment in voice-based automated services. It comes at an extraordinary time when awareness of speech recognition technology is growing in the consumer consciousness and when questions are starting to be asked about how far we can go in talking to machines," said Steve Young, Professor of Information Engineering at the University of Cambridge, UK.

Alum Appointed President of University of Potsdam

ICSI alum Oliver Günther has been appointed the president of the University of Potsdam, Germany. Günther was the second postdoctoral fellow to be hired by ICSI after its founding in 1986 and has been a strong supporter of the Institute since then. He took over as president of Potsdam on January 1.



Speech Researcher Nikki Mirghafori Selected to Attend Prestigious NAKFI Conference

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.



Networking Researcher Vern Paxson Receives 2010-2011 Jim and Donna Gray Faculty Award

The 2010-2011 Jim and Donna Gray Faculty Award has been given to Networking Group researcher Vern Paxson, who leads security efforts at ICSI. The award is given each year to a UC Berkeley Computer Science faculty member for excellence in undergraduate teaching. Previous winners include Architecture Group leader Krste Asanovic (2009-2010) and ICSI affiliate Dan Klein (2008-2009).



Networking Researcher Sylvia Ratnasamy Receives Yahoo! ACE Award

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



Sylvia Ratnasamy Receives Sloan Research Fellowship

ICSI external fellow Sylvia Ratnasamy has received a Sloan Research Fellowship. Ratnasamy, an assistant professor at UC Berkeley, works with ICSI's Networking Group. She is among 126 U.S. and Canadian researchers to receive the fellowship, which is given annually to young scientists and scholars.

ICSI to Collaborate with Microsoft on Conversational Human-Machine Interaction

Researchers at ICSI will work with Microsoft Corp. to advance the state of the art in human-computer interaction relying on speech and other modalities. The collaboration takes advantage of ICSI's history of excellence in speech processing research and Microsoft's wealth of data, technology, and experience in deploying natural speech interfaces in its services and applications.

FrameNet Director Charles Fillmore Receives ACL Lifetime Achievement Award

Professor Charles Fillmore, director of the FrameNet Project, received the Association for Computational Linguistics's Lifetime Achievement Award on July 11. The award is given for widely recognized, sustained, and enduring contributions to the field of computational linguistics over a long period. Fillmore has been and continues to be a major contributor to contemporary linguistics. The FrameNet Project is based on the theory of frame semantics, which Fillmore developed in the 1970s. Fillmore also developed case grammar theory in the 1970s, and his work in collaboration with Paul Kay and George Lakoff was generalized into the theory of construction grammar.



FrameNet Members Win Antonio Zampolli Prize

Professor Charles Fillmore and Collin Baker of the FrameNet Project are co-winners of the Antonio Zampolli Prize, which recognizes outstanding contributions to the advancement of language resources and language technology evaluation. Fillmore, the founder and director of FrameNet, and Baker, its project manager, share the prize with the Oriental Committee for the Co-Ordination and Standardisation of Speech Databases and Assessment Techniques (Oriental COCOSDA). The prize was announced on May 25 during the closing session of the Language Resources Evaluation Conference (LREC) in Istanbul, Turkey.



Research Project Seeks to Identify Underlying Challenges to Current ASR Technology

A new research project at ICSI is focused on exploring automatic speech recognition to understand the limitations and challenges from current technologies. Sponsored by the Intelligence Advanced Research Projects Activity (IARPA) via the Air Force Research Lab (AFRL), the research aims to use its conclusions to lead to new methods for improving ASR technology.

Algorithms Leader Richard Karp to Speak at Turing Award Centenaries

Professor Richard Karp, leader of the Algorithms Group, has been invited to speak at three conferences celebrating the 100th birthday of the British mathematician and computer scientist Alan Turing. Karp received the Turing Award in 1985 for his contributions to the theory of computational algorithms. The Turing Award is the most prestigious technical award given by the ACM and widely considered the most prestigious award in computing.

New Service Checks Users' SSL Certificates


Networking Group 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.

Faculty Associate Receives 2012 Goedel Prize

Faculty associate Christos Papadimitriou and his collaborator are one of three teams sharing the 2012 Gödel Prize. Papadimitriou and Elias Koutsoupias are recognized for their paper "Worst-Case Equilibria," which examines the effect of selfish networking behavior. The prize is awarded annually for outstanding papers in theoretical computer science by ACM's Special Interest Group on Algorithms and Computation Theory and the European Association for Theoretical Computer Science. It was presented on July 12 at the International Colloquium on Automata, Languages and Programming. Papadimitriou is the C. Lester Hogan Professor of Electrical Engineering and Computer Science at UC Berkeley.

Speech Researchers Win Best Poster Award

Speech Group researchers Gerald Friedland and Jaeyoung Choi won a best poster award at the Korean Student Technical and Leadership Conference. The researchers were able to automatically estimate where some YouTube and Flickr videos were shot by semantically analyzing the tags added to them by users. The researchers ran text from the tags through several filters to account for ambiguities, such as words that could refer to multiple things (for example, Paris, France and Paris, Texas) or that contained typos. The paper extends previous work showing it is possible to, for example, use the locations of videos and photos to quickly and easily find the home addresses of users currently on vacation, leaving them vulnerable to burglary.



Video Location Estimation System Receives Honorable Mention at MediaEval

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.

Alumni Receive IEEE Hamming Medal

Michael Luby, a long-time collaborator and the former leader of the Algorithms Group, received the IEEE Richard W. Hamming Medal on June 30. Luby shares the award with ICSI alum Amin Shokrollahi. They are recognized for developing efficient and flexible data coding methods that have enabled the success of information distribution applications such as video streaming and delivery of data to mobile devices. Luby is the vice president of technology at QUALCOMM's new office in Berkeley, and Shokrollahi is a professor of math and computer science at EPFL and the former chief scientist of Digital Fountain in Lausanne, Switzerland. He is also currently the CEO of Kandou Bus, a company he co-founded in 2011.

CACM Features Networking Researcher Nicholas Weaver in Roundtable Discussion

The Communications of the ACM recently featured a roundtable discussion among Networking Group researcher Nicholas Weaver and other computer scientists about bufferbloat, the overuse of buffering inside a network, which can result in slow network connections.



International Computer Science Institute
1947 Center Street, Suite 600
Berkeley, CA 94704
v: +1 510 666 2900
www.icsi.berkeley.edu

CREDITS

Design: ATELIER GRAND & CÉLINE AYMON FOURNIER (COURTESY OF IDIAP)
Photography: ALBERT PARK, ROBERTO PIERACCINI, LEAH HITCHCOCK & CHRIS SWITZER
Add'l Photography: P. 6, FLICKR USER 'VGM8383'
P. 8, FLICKR USER 'CENTRALASIAN'
P. 30, FLICKR USER 'KASEI'
P. 70, FLICKR USER 'PICTOSCRIBE'