

The ICSI GAZETTE

volume ten | issue two | may 2012

featured icsi research:
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 English, Persian, Russian, and Spanish. The work will span the fields of computer science, linguistics, psychology, and cognitive science. MetaNet, funded by Intelligence Advanced Research Projects Activity (IARPA) under its Metaphor Program, is led by Professor Srinu Narayanan, ICSI's AI Group leader.

metaphor built on more basic metaphors that underlie English speakers' cognitive processes, such as *Action Is Motion* and *Purposes Are Destinations*. These metaphors give rise to a conceptualization of love as a journey and romantic problems as obstacles along that journey. This ultimately leads to linguistic metaphors – phrases and expressions that rely on and evoke conceptual metaphors, such as “Their marriage is on the rocks.”

THEORETICAL FOUNDATIONS

The AI Group has a long and productive history 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 Professor Jerry Feldman,

Lakoff, 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.

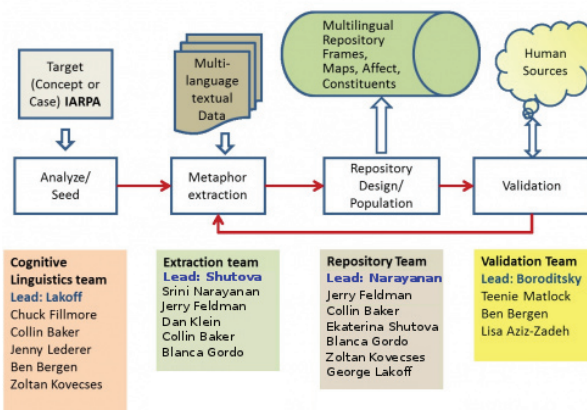
Narayanan, who in addition to leading the AI Group is a cognitive scientist on the UC

Berkeley faculty, developed the computational theory and an early interpretation system that demonstrated the utility of conceptual metaphor for reasoning about abstract concepts and domains. The neurally motivated model and results have led to new insights on the nature of inference, called the “simulation semantics” hypothesis, which asserts that understanding is imaginative simulation. Simulation semantics makes

Metaphors describe one thing in terms of another; they're mappings from a source term to a target term. The target, which is abstract, is understood in terms of the source, which is concrete. Shakespeare's “bank and shoal of time,” for example, maps abstract concepts of time and life to more concrete images having to do with bodies of water.

But metaphors are not just literary devices; they pervade language and affect, not just the way we speak about the world, but also the way we think about the world. “Many of our abstract concepts are learned through bodily experience and are understood by projecting to bodily experience,” said Narayanan, the principal investigator of the project. Conceptual metaphors – such as the idea that time is money or life is a journey – draw on these bodily experiences to express complex, abstract ideas.

Conceptual metaphors, the researchers say, structure how we see the world. UC Berkeley Professor George Lakoff pointed to the example Love Is A Journey, a complex

in this
issue:

page 2
Reflections
by Roberto Pieraccini, Director

page 3
News Briefs

page 7
Visiting Scholars

page 8
Featured Research
Twitter Spam

page 10
Publications



photo by Leide Porcu

reflections by Roberto Pieraccini, Director

I took the job as director of the International Computer Science Institute on January 15, succeeding Professor Nelson Morgan, who led the institute for 13 years. During this time ICSI went from financial uncertainty to great stability, and thrived as one of the few independent, nonprofit research centers in computer science in the U.S. Nelson Morgan will remain at ICSI as its deputy director and as the leader of the Speech Group. I will use this column as an opportunity to share my reflections on science and technology in general, the forces that shape their progress, and how they are reflected in the work at ICSI.

But first I want to say a few words about Nelson Morgan, because without him, I am sure, ICSI would not be the great place it is now, and perhaps would not exist at all. Morgan is one of those few great men you rarely meet in life whose greatness is expressed not only by their high standards of professionalism and integrity, but also their utmost humility. I have known Morgan for more than 20 years, as a colleague and fellow scientist working in speech technology. I have always admired him as one of the researchers who follow the less trodden path with the goal of making a real difference in science. Indeed he has. I have known him mostly for his original groundbreaking work on speech processing, but as I started working with him over the past few months I realized how much he has done for ICSI, and how much of his time he has dedicated to it. That certainly sets a high bar for me as the new director.

In these first few months at ICSI I have asked everyone the same question, "What is ICSI for you," and among many answers, one was pretty constant, especially among old timers: ICSI is my home. Some told me that after finishing their dissertation work at ICSI and receiving their doctorate, they went to work in industry for several years before they decided to return to ICSI. And the feeling was unmistakable: "I am back home!" ICSI is home to many world-class researchers, whether they

are still here or not, and I want this to remain the case for them and for future generations of computer scientists and technologists. There is no place like home, there is no place like ICSI.

But besides being a home for many great scientists, ICSI is a place where science is pursued, technology created, new industries foreshadowed. There is always something new, something great going on here, from ideas about the future of the Internet to the discovery of sophisticated network security breaches that may affect all of society, from the modeling of complex and not-yet-understood linguistic phenomena, to attempts to shed light on the inner workings of the brain. Our scientists are recreating the complex mechanisms that allow humans to understand speech in spite of noise, accents, and linguistic variations; they are showing how to use the wealth of multimedia Web data to find additional related information and how to prevent its use in crime. They are modeling the complex relationships between vision and language acquisition, and they are discovering effective algorithms to help understand the complex relations among elements of genetics. Some of our scientists focus on advancing the theory of computer science and algorithms, and others on realizing computers based on new physical principles. This is ICSI, a relatively small institute with about a hundred people between researchers, visiting scholars, and staff, independent yet located next to UC Berkeley, one of the most exciting centers of computer science and engineering research in the world.

Now the question for me is how I can help all of this greatness to continue on its path and to grow. ICSI is strong and will continue to grow stronger; that is my mission. We will continue to attract the most talented young researchers, strengthen the relationships with our partner countries and establish new ones, identify potential areas of research and

seed new teams, work with industry to invent a new future that will improve our society. I still have many questions about how to achieve all of that, but I am sure I will find some answers with the help of everyone here at ICSI.



news briefs

ROBERTO PIERACCINI is ICSI's new director. A renowned leader in science and technology, Pieraccini brings expertise in both academic and corporate research as well as enterprise business leadership. Before joining ICSI, Pieraccini was the chief technology officer of SpeechCycle, a company focused on advanced speech and multimodal solutions, and was instrumental to the most recent advancements in spoken dialog interaction technology. He has also been a research staff member at IBM's T.J. Watson Research and a director of research and development for SpeechWorks. He has served as the chair of IEEE's Speech and Language Technical Committee and has worked for AT&T Laboratories and Bell Labs. He started his career as a researcher at CSELT, the research labs for the Italian telephone company in the 1980s. He serves on the editorial boards for several science and technology magazines and is a fellow of IEEE and the International Speech Communication Association. Pieraccini succeeds Professor Nelson Morgan, who served as director for 13 years and helped found the Speech Group. Morgan will now act as deputy director and continue to lead the Speech Group.

Professor **RICHARD KARP**, leader of the 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 like social sciences, biology, physics, and economics.

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.

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

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. The prize was announced on May 25 during the closing session of the Language Resources Evaluation Conference in Istanbul, Turkey.

Director **ROBERTO PIERACCINI**'s book *The Voice in the Machine*, about the evolution of speech recognition technology, was released March 30 by MIT Press. The book examines six decades of work in science and technology to develop computers that can interact with humans using speech and the industry that has arisen around these technologies.



Networking Group researcher **BARATH RAGHAVAN** and UC Berkeley postdoctoral fellow Justin Ma have found that the Internet and devices that access it are responsible for about 2 percent of global energy usage. The researchers estimated how much energy it takes to manufacture end systems, such as laptops and smart phones, and how much energy these devices use. They estimated that the Internet uses between 170 and 307 gigawatts of the 16 terawatts used globally. The findings, presented at HotNets, suggest that research into energy efficiency should focus less on reducing the power used by networking and more on ways that networking can reduce other forms of energy use, for example, by replacing business travel with teleconferencing.

OMOJU MILLER, a graduate student researcher in the AI Group, attended two events to discuss social justice and technology, a focus of her work. She was invited to participate in the inaugural TEDx Summit from April 16 to 20 in Qatar. The TEDx program sponsors local, independent events licensed by TED, the popular series of conferences on technology, entertainment, and design. Last year, Miller organized TEDxEuclidAve, which focused on technology and entrepreneurship used to achieve social good. She also participated in a panel discussion on social justice and new media technologies as part of a conference on inclusion and diversity in the age of data, held April 26.

continued on page 9

featured research: continued

continued from page 1

detailed predictions on the manner, content, and timing of metaphoric inference. These predictions continue to be confirmed by experimental results in laboratories around the world. Narayanan's more recent work demonstrates through computer simulation how metaphors emerge as a natural consequence of the functional architecture and known information processing constraints of the human brain.

"The knowledge model and representation techniques in the existing system will significantly inform the design and implementation of the metaphor repository and its use in applications of interest to IARPA in the program," Narayanan said. "This work enables us to hit the ground running."

The team's linguistic analysis lead, George Lakoff, is a linguist, cognitive scientist, and co-founder of the NTL project at ICSI and UC Berkeley. Over the last three decades, Lakoff and his students have pioneered much of the modern research on metaphor. This research, reported in more than a dozen books, shows that metaphoric linguistic expressions are surface manifestations of metaphorical thought processes, and that a large class of conceptual metaphors arise from bodily experiences – which began the research tradition of what is now known as "embodied cognition." Lakoff and Feldman founded the NTL project, an integral part of the AI Group since 1988.

In recent analyses of metaphors used in political discourse, Lakoff has uncovered basic parenting frames and foundational metaphor systems that underlie deeply held beliefs and worldviews. These metaphor systems are of core interest and will inform the linguistic research in the project.

A MULTI-LINGUAL REPOSITORY

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 on metaphor) from text and interprets them automatically in four different languages. Researchers in ICSI's MetaNet project will do 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. While

different languages use different metaphors, the use of metaphor "appears to be cross-linguistic," Narayanan said. "You and I, from different backgrounds, share this."

For example, speakers from all cultures studied for use of metaphor discuss abstract ideas, such as political action, in terms of spatial motion, such as physical movement. The metaphor underlying this phenomenon – the event structure metaphor – is so basic and intuitive that it shapes the way we think about abstract actions without our knowing. UC Merced Professor Teenie Matlock, an experimental cognitive linguist involved in the validation portion of the project, says the way we speak about time is an example of this. English speakers think of time in spatial terms: April is *after* March. This metaphor influences not just the words and phrases we use to discuss time, but also the way we conceive of time. Matlock said, "We often go to the physical sense of time to help us think about it."

IARPA, a research arm of the Office of the Director of National Intelligence (ODNI), hopes to use the results of the Metaphor Program to understand the role metaphor plays in how people from different cultural backgrounds make judgments and decisions. The system developed during the MetaNet project will help by automatically analyzing metaphor in the four languages under investigation. The project, said Narayanan, will "try to understand how different cultures have different beliefs and worldviews as a result of these metaphors."

Automatic text analyzers have difficulty dealing with metaphor. One reason is that the metaphorical mapping from a concrete source to an abstract target is only partial. For example, while *information*, an abstract target, can be discussed in terms of a liquid – it may be contained or leaked – the metaphor does not extend to every aspect of a liquid.

Another challenge is that complex metaphors arise from more basic metaphors. As Lakoff pointed out, the metaphors *Action Is Motion* and *Goals Are Destinations*, when combined with others, eventually lead to the complex idea that *Love Is A Journey*. A computer system must understand the relationships among these metaphors before it can understand that a marriage "on the rocks" is a troubled one.

IARPA is providing the team with a series of target concepts to be analyzed by the system in the first years of the project; later, the system will be asked to respond to test cases and scenarios determined by IARPA.

FRAME SEMANTICS

The analysis of target concepts and cases will rely in part on frame semantics work done by ICSI's FrameNet Project. FrameNet, led by Collin Baker and Professor Charles Fillmore, is building a lexical database, usable by both machines and humans, that shows how words are used in texts. In a sentence annotated by FrameNet, a word is understood through the semantic frame it evokes and the roles it plays in that frame. A semantic frame is a description of a type of event, relation, or entity, and comprises frame elements that may include, for example, the person carrying out an action and the object on which it is carried out.

FrameNet annotations go beyond grammatical parsing. The agent of a frame – the one carrying out an action – may not necessarily be the subject of a sentence that evokes the frame. This is common in commands (“Leave that alone”), in which the agent you is implied, but is also true, for example, in the sentence “Graham’s back arched,” where the agent is Graham even though the grammatical subject is Graham’s back.

This method of annotation will be used to describe the relationship between a linguistic metaphor’s literal meaning and its metaphorical meaning. For a word that evokes a metaphor, cognitive linguists will determine both the target and the source. In the sentence “That flat tire cost me an hour,” the annotation of the word cost will show its abstract target (*time*) and its concrete source (*money*). This will provide the repository with information about how metaphorical frames and frame elements work together, forming a network that shows the relationships between metaphors and their literal meanings. The repository will also show the relationships among different metaphors.

These mappings will be used as base information by a system that extracts metaphors using supervised machine learning. The system will use what it has learned about the relationships between the frame elements of conceptual metaphors to find new metaphors in text.

AUTOMATIC EXTRACTION

The team is also working toward unsupervised machine learning of metaphor mappings, which uses less initial information than supervised learning and so will be particularly helpful for extracting metaphors in languages, such as Persian, with fewer resources. While a graduate student at the University

of Cambridge, Ekaterina Shutova, now an ICSI postdoctoral fellow, developed a system able to extract metaphors using a small set of annotated metaphors. Shutova clustered words associated with these metaphors, and the system used these word clusters to find new metaphors that used the same words. Shutova points to the metaphor *Relationship Is A Mechanism* as an example. If her system were seeded with this metaphor and a cluster that grouped the words relationship and function, it would be able to identify the metaphor *Democracy Is A Mechanism*, since democracy can be seen in similar linguistic contexts in the data. While the system operates with high precision, it is dependent on its seed metaphors, which currently results in low recall. Shutova’s next goal is to expand on this work and build a more robust, potentially fully unsupervised system. With these extraction methods and the representation techniques, the researchers will fill a repository of metaphors and their frames, mappings, and elements.

EVALUATING THE SYSTEM

In conjunction with this, a group of cognitive linguists and neuroscientists will test how metaphor affects thinking and emotion in order to evaluate the effectiveness of the metaphor repository and methodologies developed by the rest of the team. The validation team is led by Lera Boroditsky of Stanford and includes Lisa Aziz-Zadeh of the University of Southern California, Ben Bergen of UC San Diego, and Teenie Matlock of UC Merced

Matlock, founder of the Cognitive and Information Science program at UC Merced, says this part of the project will “test the underlying psychological reality of metaphor.”

“We want to know what are the conceptual underpinnings of metaphor in actual conversation,” she said.

Although the validation work is in the early stages of planning, the team hopes to use, among other techniques, eye-tracking tests, in which participants look at a blank screen and listen to an audio track while a device records the movement of their eyes. Such tests reveal changes in participants’ internal state. There will be similar analyses of participants’ gestures and posture while they engage in conversation.

Matlock specializes in conducting tests using natural discourse, in which participants watch a video or read a paragraph and either answer questions about it or discuss it in pairs. Researchers analyze the transcribed recording of the

continued from page 5

discussions, tracking what words are used, in what order they are used, and how frequently they are used.

The team will also evaluate how effectively the system extracts metaphors by asking study participants to rate the metaphors on a variety of measures, including how useful and familiar they are.

The team is also interested in “seeing whether people have implicit emotional reactions to the various metaphors,” said Bergen, a behavioral researcher at UC San Diego. He will be running implicit association tests (IATs) on the metaphors extracted by the system. IATs measure unconscious reactions by measuring how long it takes participants to categorize terms along a set of axes, one of which has a value judgment associated with it. IATs are often used to measure racial bias: participants are asked to categorize faces that appear on a screen as one of two races, which are paired with value judgments (such as “good” and “bad”). Participants who quickly categorize minority faces as the correct race both when the value judgment is positive and when it is negative are said to have low racial bias. Bergen will adopt this experimental design, substituting metaphors for faces.

The team will also look at biophysical responses to the use of metaphor. Bergen will use electromyographic recordings to measure subtle facial expressions that may be so slight that they are not visible. Aziz-Zadeh of the University of Southern California will track what parts of the brain are activated by metaphor using functional Magnetic Resonance Imaging.

The results of these tests will give researchers an idea of how metaphors affect speakers’ brain activity and emotion and how they influence decision-making and judgment. The results will also inform the work of the rest of the team as they develop the metaphor extraction system and methodology.

Baker, the FrameNet and MetaNet project manager, says that the combination of disciplines makes the project unique. “Our approach is particularly challenging because it requires bridging vastly different fields, from cognitive linguistic analysis to machine learning techniques to brain scans,” he said. “But if we can do it, it will be particularly rewarding.”

THE IMPLICATIONS OF METAPHOR

Previous brain imaging studies have shown that talking about metaphorically grasping an idea uses the same parts of the

brain as physically grasping an object. This strong connection between metaphor and bodily experience, Narayanan said, has implications for political and social discourse.

Narayanan pointed out that images of strangulation are often used in political discourse (“our economy is being strangled by socialist policies”). Narayanan said, “You’re using a word that has a very deep connection to our experience as humans to talk about government.”

“Language understanding is imaginative simulation,” he said.

Boroditsky, the Stanford professor leading the experimental validation portion of the project, and her student Paul Thibodeau found in previous studies that changing which metaphors are used to discuss a crime wave affects how people think it should be solved. When crime was described as a beast “lurking in neighborhoods,” participants in a study were more likely to suggest increased enforcement – more police officers and harsher punishments – as a solution. When the same crime wave was described as an illness “plaguing” neighborhoods, participants were more likely to suggest social reform.

Narayanan said this raises several questions, such as, “Why these mappings and not others? Are there different sets of mappings across cultures? Can we intervene and suggest different mappings to change the way people think about crime?”

Matlock has done similar work showing that the way something is framed affects how people think about it. In one test, some participants read a paragraph about a political candidate who “was having an affair”; others read about a candidate who “had an affair.” The former believed that the candidate was less likely to be elected than the latter. Such findings, said Matlock, suggest that our thinking is influenced, not just by the use of words and metaphors, but by the form of the linguistic construction as well.

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. 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, DoD/ARL, or the U.S. Government.

visiting scholars

september 2011 – march 2012

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

AI

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

ALGORITHMS

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

ARCHITECTURE

Miquel Moretó Planas

CAMPUS AFFILIATION / OTHER

Nils Peters (Germany)

NETWORKING

Bernhard Amann (Germany)
Mohan Dhawan
Haixin Duan
Oana Goga
Dorgival Guedes (Brazil)
Junaid Khalid
Emil Lagerspetz (Finland)
Gregor Maier (Germany)
Syed Akbar Mehdi
Sasu Tarkoma (Finland)
Renata Teixeira
Amin Tootoonchian
Jack Wu
Andreas Wundsam (Germany)

SPEECH

Seppo Enarvi (Finland)
Arlo Faria (Brazil)
Mikko Kurimo (Finland)
Bernd Meyer (Germany)

VISION


Tim Althoff
Tobias Baumgartner
Matthias Kirchner (Germany)
Arto Klami (Finland)
Seppo Virtanen (Finland)



Seppo Virtanen



Miguel Moretó



Amin Tootoonchian



Gerard de Melo

featured research: twitter spam

Spammers who posted almost half a million Twitter messages in order to silence debate over Russia's election in December 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 Group 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, 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.

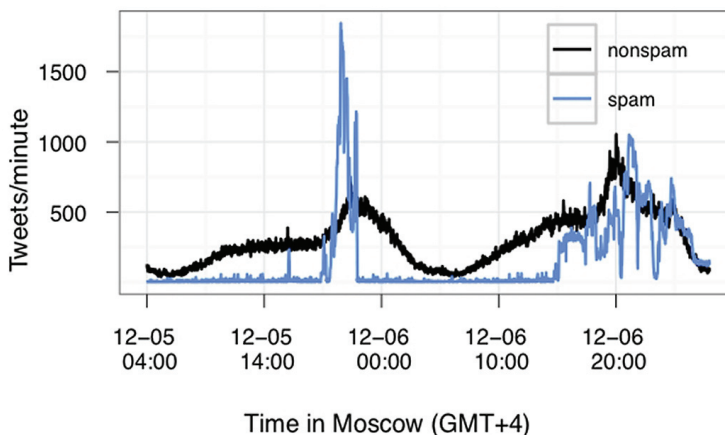
In addition, the researchers also found that the IP addresses used to post spam about the election were far more dispersed around the globe than those used to post to legitimate messages, which tended to originate in Russia. The spam IP addresses were also more likely to appear on the Composite Blocking 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.



news briefs:

continued

Speech Group researcher **GERALD FRIEDLAND** is the 2011 associate editor of the year of the *ACM Transactions on Multimedia Computing, Communications, and Applications*. Friedland was recognized for excellent performance as an associate editor.

Lifehacker.com has featured Priv3, an extension for Firefox developed by **NETWORKING GROUP** members that keeps social networking sites from tracking users' movement on certain Web pages without their knowledge.

Networking Group Leader **SCOTT SHENKER** was elected to the National Academy of Engineers, one of the highest distinctions accorded to engineers. Shenker was recognized for 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. Other members of the National Academy of Engineering include Algorithms Group leader Richard Karp and ICSI Board of Trustees members David Culler and Prabhakar Raghavan.

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).

ICSI held its annual BEARS Open House February 23. Professor **SRINI NARAYANAN**, leader of the AI Group, led a panel discussion on MetaNet with Professor **GEORGE LAKOFF** and **EKATERINA SHUTOVA**. Scientists from all groups presented posters summarizing recent results of various research projects. The Open House is held annually in conjunction with UC Berkeley EECS Annual Research Symposium (BEARS).



Pieraccini and Steve Wegmann at BEARS Open House.

Alum **MARK HANDLEY** has received the 2012 IEEE Internet Award. Handley, now a professor at University College London, was recognized for his contributions to Internet multicast, telephony, congestion control, and the shaping of open Internet standards and open-source systems in all these areas. Past winners of the award include Networking Group Leader Scott Shenker in 2006 and Networking Group researcher Sally Floyd in 2005.

AI Group Leader **SRINI NARAYANAN** and his collaborators have received the Semantic Web Science Association's Ten-Year Award. The award recognizes papers that have had the most influence a decade after being presented at the International Semantic Web Conference. This is the first year the award has been given.

AI Group researchers, with a colleague at UC Irvine, are studying how Pacific Rim cultures categorize colors differently and creating a public database that shows how speakers of 116 Mesoamerican languages name different colors. The Mesoamerican Color Survey, collected between 1978 and 1981 by the late cognitive anthropologist Robert E. MacLaury, documents how cultures in Mexico and Central America categorize color. Until now, the data has been available only in hard copy. Kimberly A. Jameson, associate project scientist at UC Irvine's Institute for Mathematical Behavioral Sciences, is working with AI Group researchers **PAUL KAY** and **RICHARD S. COOK** to digitize the data. The data will be hosted at ICSI alongside the World Color Survey, a digital archive of data showing how speakers of 110 unwritten languages spoken around the world categorize color. The project is supported in part by the UC Pacific Rim Research Program.

Congratulations to Networking Group senior researcher **MARK ALLMAN** and his wife Meredith on the birth of their son, Ryan Thomas, who was born on October 26 at 2:00 p.m. Ryan Thomas weighed 8 pounds 14 ounces and was 20.5 inches at birth.



Ryan Thomas Allman

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.

O. BELL, M. ALLMAN, AND B. KUPERMAN. On Browser-Level Event Logging. ICSI Technical Report TR-12-001, January 2012.

T. CALLAHAN, M. ALLMAN, AND M. RABINOVICH. Psst, Over Here: Communicating Without Fixed Infrastructure. ICSI Technical Report TR-12-002, January 2012.

J. CHOI AND G. FRIEDLAND. Data-Driven vs. Semantic-Technology-Driven Tag-Based Video Location Estimation. *Proceedings of the IEEE International Conference on Semantic Computing (ICSC 2011)*, Palo Alto, California, pp. 243-246, September 2011.

J. CHOI, H. LEI, AND G. FRIEDLAND. The 2011 ICSI Video Location Estimation System. *Proceedings of the MediaEval 2011 Workshop*, Pisa, Italy, September 2011.

L. CONDE, S. BEVAN, M. SITZER, N. KLOPP, T. ILLIG, J. THIERY, J. SESSLER, J. BAUMERT, O. RAITAKARI, M. KHINEN, L. P. LYYTIKINEN, R. LAAKSONEN, J. VIKARI, T. LEHTIMKI, W. KOERNIG, E. HALPERIN, AND H. S. MARKUS. Novel Associations for Coronary Artery Disease Derived from Genome Wide Association Studies Are Not Associated with Increased Carotid Intima-Media Thickness.... *Atherosclerosis*, Vol. 219, No 2, pp. 684-689, December 2011.

O. CULO, S. HANSEN-SCHIRRA, K. MAKSYMSKI, AND S. NEUMANN. Empty Links and Crossing Lines: Querying Multilayer Annotation and Alignment in Parallel Corpora. In *Translation: Computation, Corpora, Cognition (TC3)*, Special Issue on Parallel Corpora: Annotation, Exploitation, Evaluation, Vol. 1, No. 1, pp. 75-104, December 2011.

M. ELLSWORTH. Toward a Representation for Understanding the Structure of Multiparty Conversations. ICSI Technical Report TR-11-003, November 2011.

R. FARRELL, O. OZA, N. ZHANG, V. MORARIU, T. DARRELL, AND L. DAVIS. Birdlets: Subordinate Categorization Using Volumetric Primitives and Pose-Normalized Appearance. *Proceedings of the International Conference on Computer Vision (ICCV)*, Barcelona, Spain, pp. 161-168, November 2011.

G. FRIEDLAND. Using a GPU, Online Diarization = Offline Diarization. ICSI Technical Report TR-12-004, January 2012.

G. FRIEDLAND, J. CHOI, AND A. JANIN. Video2GPS: A Demo of Multimodal Location Estimation on Flickr Videos. *Proceedings of the ACM Multimedia Conference (MM'11)*, Scottsdale, Arizona, November 2011.

G. FRIEDLAND, J. CHOI, H. LEI, AND A. JANIN. Multimodal Location Estimation on Flickr Videos. *Proceedings of the ACM International Workshop on Social Media (WSM11)*, Scottsdale, Arizona, November 2011.

G. FRIEDLAND, J. FISCUS, T. HAIN, AND S. FURUI, EDS. Special Section on New Frontiers in Rich Transcription. *IEEE Transactions on Audio, Speech, and Language Processing*, Vol. 20, No. 2, February 2012.

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, G. MAIER, R. SOMMER, AND N. WEAVER. Sherlock Holmes's Evil Twin: On The Impact of Global Inference for Online Privacy. *Proceedings of the New Security Paradigms Workshop (NSPW)*, Marin County, California, September 2011.

E. FRIEDMAN. Asymmetric Cost Sharing Mechanisms. *Games and Economic Behavior*, available at <http://dx.doi.org/10.1016/j.geb.2011.09.009>, October 2011.

A. GHODSI, T. KOPONEN, B. RAGHAVAN, S. SHENKER, A. SINGLA, AND J. WILCOX. Information-Centric Networking: Seeing the Forest for the Trees. *Proceedings of the 10th ACM Workshop on Hot Topics in Networks (HotNets-X)*, Cambridge, Massachusetts, November 2011.

A. GHODSI, T. KOPONEN, B. RAGHAVAN, S. SHENKER, A. SINGLA, AND J. WILCOX. Intelligent Design Enables Architectural Evolution. *Proceedings of the 10th ACM Workshop on Hot Topics in Networks (HotNets-X)*, Cambridge, Massachusetts, November 2011.

D. GILLICK, L. GILLICK, AND S. WEGMANN. Don't Multiply Lightly: Quantifying Problems with the Acoustic Model Assumptions in Speech Recognition. *Proceedings of the Automatic Speech Recognition and Understanding Workshop (ASRU)*, Big Island, Hawaii, December 2011.

B. GOLD, N. MORGAN, AND D. ELLIS. *Speech and Audio Signal Processing: Processing and Perception of Speech and Music*, 2nd Edition. Wiley, 2011.

E. GONINA, G. FRIEDLAND, H. COOK, AND K. KEUTZER. Fast Speaker Diarization Using a High-Level Scripting Language. *Proceedings of the IEEE Automatic Speech Recognition and Understanding Workshop (ASRU 2011)*, Big Island, Hawaii, December 2011.

Y. HASEGAWA, R. LEE-GOLDMAN, A. KONG, K. OHARA, K. AKITA, S. UCHIDA, AND C. J. FILLMORE. A Study of Paraphrasing in Terms of FrameNet and Frame Semantics. *Proceedings of the 6th International Conference on Construction Grammar (ICCG-6)*, Prague, Czech Republic, September 2011.

L. HODGKINSON AND R. M. KARP. Algorithms to Detect Multiprotein Modularity Conserved During Evolution. *IEEE/ACM Transactions on Computational Biology and Bioinformatics*, PMID: 21968956, September 2011.

F. HOPFGARTNER AND J. M. JOSE. Development of a Test Collection for Studying Long-Term User Modelling. *Proceedings of the Workshop Information Retrieval (LWA'11)*, Magdeburg, Germany, September 2011.

A. JANOCH, S. KARAYEV, Y. JIA, J. Y. BARRON, M. FRITZ, K. SAENKO, AND T. DARRELL. A Category-Level 3-D Object Dataset: Putting the Kinect to Work. *Proceedings of the First IEEE Workshop on Consumer Depth Cameras for Computer Vision at the International Conference on Computer Vision (ICCV)*, Barcelona, Spain, pp. 1168-1174, November 2011.

Y. JIA AND T. DARRELL. Heavy-Tailed Distances for Gradient Based Image Descriptors. Proceedings of the 25th Annual Conference on Neural Information Processing Systems (NIPS 2011), Granada, Spain, December 2011.

C. KREIBICH, N. WEAVER, C. KANICH, W. CUI, AND V. PAXSON. GQ: Practical Containment for Measuring Modern Malware Systems. Proceedings of the 2011 Internet Measurement Conference (IMC), Berlin, Germany. Also appeared as ICSI Technical Report TR-11-002., November 2011.

Y. LEE. Efficient VLSI Implementations of Vector-Thread Architectures. UC Berkeley Master's thesis, Berkeley, California. Also appeared as EECS Department, UC Berkeley Technical Report No. UCB/EECS-2011-129, Berkeley, California, December 2011.

J. LIU, S. SHENKER, M. SCHAPIRA, AND B. YANG. Data-Driven Network Connectivity. Proceedings of the 10th ACM Workshop on Hot Topics in Networks (HotNets-X), Cambridge, Massachusetts, November 2011.

C. LUQUE, M. MORETÓ, F. J. CAZORLA, R. GIOIOSA, A. BUYUKTOSUNOGLU, AND M. VALERO. CPU Accounting for Multicore Processors. IEEE Transactions on Computers. Vol. 61, Issue 2, February 2012.

A. MARON-KATZ, E. B. SIMON, Y. JACOB, K. ROSENBERG, R. M. KARP, T. HENDLER, AND R. SHAMIR. Using Contiguous Bi-Clustering for Data Driven Temporal Analysis of fMRI Based Functional Connectivity. Proceedings of the 4th Annual INCF Neuroinformatics Congress, Boston, Massachusetts, September 2011.

R. MERTENS, P.-S. HUANG, L. GOTTLIEB, G. FRIEDLAND, AND A. DIVAKARAN. On the Applicability of Speaker Diarization to Audio Concept Detection for Multimedia Retrieval. Proceedings of the IEEE International Symposium on Multimedia, Dana Point, California, pp. 446-451, December 2011.

R. MERTENS, H. LEI, L. GOTTLIEB, G. FRIEDLAND, AND A. DIVAKARAN. Acoustic Super Models for Large Scale Video Event Detection. Proceedings of the ACM International Workshop on Events in Multimedia (EiMM11), Scottsdale, Arizona, November 2011.

B. T. MEYER. Improving Automatic Speech Recognition by Learning from Human Errors. Proceedings of the 162nd Meeting of the Acoustical Society of America, San Diego, California, October 2011.

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.

R. MORRISON, E. FRIEDMAN, AND A. LANDSBERG. Combinatorial Games with a Pass: A Dynamical Systems Approach. Chaos, Vol. 21, No. 4, December 2011.

B. RAGHAVAN AND J. MA. The Energy and Emery of the Internet. Proceedings of the 10th ACM Workshop on Hot Topics in Networks (HotNets-X), Cambridge, Massachusetts, November 2011.

M. RIBEIRO, P. GUERRA, L. VILELA, A. VELOSO, D. GUEDES, W. MEIRA JR., M. CHAVES, K. STEDING-JESSEN, AND C. HOEPERS. Spam Detection Using Web Page Content: A New Battleground. Proceedings of the 8th Annual Collaboration, Electronic Messaging, Anti-Abuse and Spam Conference (CEAS 2011), Perth, Australia, September 2011.

K. SAENKO, S. KARAYEV, Y. JIA, A. SHYR, A. JANOSH, J. LONG, M. FRITZ, AND T. DARRELL. Practical 3-D Object Detection Using Category and Instance-Level Appearance Models. Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), Algarve, Portugal, pp. 793-800, September 2011.

U. SCHMID AND E. KITZELMANN. Inductive Rule Learning on the Knowledge Level. Cognitive Systems Research, Special Issue on Complex Cognition, Vol. 12, Issues 3-4, pp. 237-248, September 2011.

R. SOMMER, L. DE CARLI, N. KOTHARI, M. VALLENTIN, AND V. PAXSON. HILTI: An Abstract Execution Environment for Concurrent, Stateful Network Traffic Analysis. ICSI Technical Report TR-12-003, January 2012.

H. O. SONG, M. FRITZ, C. GU, AND T. DARRELL. Visual Grasp Affordances From Appearance-Based Cues. Proceedings of the First IEEE Workshop on Challenges and Opportunities in Robot Perception at the International Conference on Computer Vision (ICCV), Barcelona, Spain, pp. 998-1005, November 2011.

K. STEPHENS AND B. KIRKPATRICK. Efficiently Solvable Perfect Phylogeny Problems on Binary and k-State Data with Missing Values. Proceedings of the 11th Workshop on Algorithms in Bioinformatics (WABI 2011), Saarbrücken, Germany, September 2011.

L. STOLL. Finding Difficult Speakers in Automatic Speaker Recognition. UC Berkeley PhD thesis, Berkeley, California, December 2011.

K. THOMAS, C. GRIER, V. PAXSON, AND D. SONG. Suspended Accounts in Retrospect: An Analysis of Twitter Spam. Proceedings of the 2011 Internet Measurement Conference (IMC), Berlin, Germany, November 2011.

T. TUYTELAARS, M. FRITZ, K. SAENKO, AND T. DARRELL. The NBNB Kernel. Proceedings of the International Conference on Computer Vision (ICCV), Barcelona, Spain, pp. 1824-1831, November 2011.

J. WANG, G. J. GEESMAN, S. L. HOSTIKKA, M. ATALLAH, B. BLACKWELL, E. LEE, P. J. COOK, B. PASANIUC, G. SHARIAT, E. HALPERIN, M. DOBKE, M. G. ROSENFELD, I. K. JORDAN, AND V. V. LUNYAK. Inhibition of Activated Pericentromeric SINE/Alu Repeat Transcription in Senescent Human Adult Stem Cells Reinstates Self-Renewal. Cell Cycle, Vol. 10, No. 17, pp. 3016-3030, September 2011.

P. C. WANG, S. MILLER, M. FRITZ, T. DARRELL, AND P. ABBEEL. Perception for the Manipulation of Socks. Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), Algarve, Portugal, pp. 4877-4884, September 2011.

J. YANGQING, M. SALZMANN, AND T. DARRELL. Learning Cross-Modality Similarity for Multinomial Data. Proceedings of the International Conference on Computer Vision (ICCV), Barcelona, Spain, pp. 2407-2414, November 2011.

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.

M. ZAHARIA, W. J. BOLOSKEY, K. CURTIS, A. FOX, D. PATTERSON, S. SHENKER, I. STOICA, R. M. KARP, AND T. SITTLER. Faster and More Accurate Sequence Alignment with SNAP. arXiv.org under arXiv:1111.5572v1, November 2011.



INTERNATIONAL
COMPUTER SCIENCE
INSTITUTE

THE INTERNATIONAL COMPUTER SCIENCE INSTITUTE (ICSI)

is one of the few independent, non-profit US labs solely conducting open, non-proprietary, pre-competitive research in computer science. ICSI's mission is to further research in computer science through international collaboration, and further international collaboration through research in computer science. Affiliated with the University of California campus in Berkeley, ICSI provides a haven for computer scientists to conduct long-term research without commercial limitations and with few faculty pressures. ICSI has significant efforts in these major research areas: Internet research, including Internet architecture, related theoretical questions, and network services and applications; theoretical computer science, including applications to bioinformatics; artificial intelligence, particularly for applications to natural language understanding; natural speech processing; computer vision; and computer architecture.

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

BOARD OF TRUSTEES

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

EDITOR-IN-CHIEF

Maria Quintana

SENIOR EDITOR

Roberto Pieraccini

DESIGN

envision ID

COPYWRITER

Chris Switzer

LAYOUT

Jacob Wolkenhauer

PHOTOGRAPHY

Jaci Considine
Leah Hitchcock

INTERNATIONAL
COMPUTER
SCIENCE
INSTITUTE

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