the newsletter of the International Computer Science Institute

The ICSI GAZETTE

volume seven | issue two | march 2009

featured research: computer vision

What kind of ICSI research could possibly tie together household service robots, wartime military operations, and an afternoon trip to IKEA? Trevor Darrell's newly formed Vision group does just that. Their work on computer vision enables recognition of objects, gestures, and even faces. As computer vision becomes more robust, it's finding applications identifying glassware, aiding in military planning, and even reuniting parents with their children.

Professor Trevor Darrell relocated from the CSAIL group at MIT and began working at ICSI in the summer of 2008. Joining him upon arrival were Dr. Raquel Urtasun and Mario Christoudias. In less than a year, the group has already grown to eight staff and student researchers here in Berkeley, also including Ashley Eden, Alex Shyr, Mario Fritz, Brian Kulis, Mathieu Salzmann, and Carl Ek. Meanwhile, Darrell is still advising his Ph.D. students who remain at MIT, Ariadna Quattoni, Tom Yeh, and Kate Saenko.

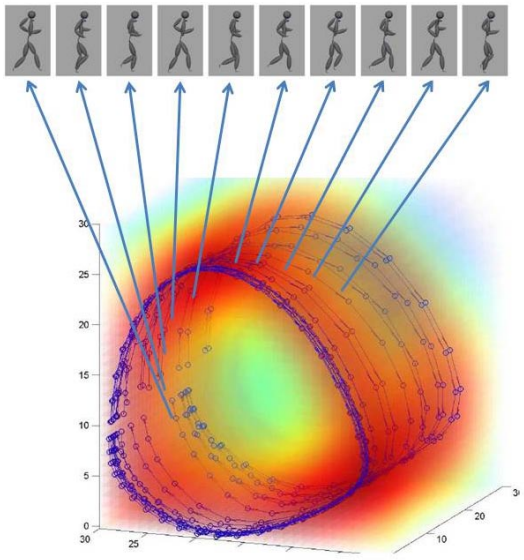

A visual representation of human body motion tracking work done by Dr. Raquel Urtasun, which can recognize and characterize a person's motion.

Darrell's group nominally revives Vision work that was done at ICSI in the 80's and 90's with Steven Omohundro and Chris Bregler, but is substantially different in both scope and function. The arrival of Darrell and his group to ICSI is a welcome elaboration on our other work on human-machine interfaces in the Speech and AI research groups. In addition to work done in poly-semantic word disambiguation, interactive image matching, and other fields, we highlight a couple of the group's current endeavors below.

REUNITE

Just before making the move to Berkeley, Trevor Darrell's group was approached by Children's Hospital Boston. They wanted his help finding a program that would assist them in disaster relief. The program they wanted would be able to sort through pictures of children who have been rescued from an emergency, and quickly identify them. The bad news was that the kind of technology they wanted didn't exist yet; the good news is that Darrell was able to partner with them and secure funding to develop it.

Ashley Eden and colleagues in the Vision group are now working on a system that can compile and analyze photos of children's faces to help reunite them with their parents. They hope to develop a system in which a user looks at two pictures and selects the one that looks more like the child they're trying to identify, thereby quickly homing in on the picture (and thus, location) of their own child. In order to do this, the group is determining which meaningful facial features are useful in

Continued on Page 4

in this issue:

as i see it by Nelson Morgan, Director

We have met the enemy, and he is us.

-- Pogo, a.k.a. Walt Kelly

I'm writing this column in January at the start of the Obama era. High hopes are all around, and at the same time great trepidation – the economy has taken a nosedive, worse than the recession that came after the dotcom bust and the 9/11 attacks. Unemployment is rising throughout the country (and the world), the stock market is in the doldrums, and there are multiple wars in the Mideast. Upcoming baby boomer retirements are set to bust the entitlement programs, while the health care system manages to be utterly insufficient and yet incredibly expensive. In the meanwhile, humanity's carbon imprint continues to grow (though a momentary slowing in its growth may be the only positive aspect of the worldwide recession) and our dependence on fossil fuels remains nearly complete. With broken budgets in California and many other states, educational systems are coming under even greater stress, with likely increases in class sizes, college tuition, and reductions in services. And however much we might point to particular politicians for having brought us to this juncture (and they have), the American people (and more generally, people around the world) are ultimately responsible. We are the enemy.

Not exactly an upbeat scenario, despite the enthusiasm with which much of the country has greeted the start of our new President's term. Clearly the new administration has had its hands full from its first day in office. However gifted the new President and his advisors might be, they are only human.

But if we are the enemy, we are also our own cavalry. "We are the ones we have been waiting for" is a nice turn of phrase, but indeed it is only the people themselves who can provide the upturn we require. We are the ones who must provide support for new programs to improve education, to toughen pollution standards, and to radically transform the health care system. We are the ones who must insist on strict conformance to constitutional rights, and on the use of science-based decision-making. It is the support of the people for all of these things that will permit the new government to move us in a better direction. Statements made by this administration, such as the ones at www.whitehouse.gov, or in previous memos released during the campaign and transition, give us hope that they will be receptive to these messages.

What does any of this have to do with ICSI? ICSI, like any other organization devoted to research and education, is deeply affected by the world at large. Trends over the last 30 years have made it increasingly difficult to pursue long-term research, as short-term goals have come to dominate. An upswing in respect for science will make it more likely to at least partly return to an era where a longer-term perspective in scientific inquiry will be supported. A radical improvement in America's image should increase the pool of young scientists who will be excited to come to the States, and in particular to Berkeley. And as ICSI is an international institute, this will be particularly important to us.

I can't talk about the effects of the new administration without at least acknowledging its historic character. Race should have no part in the choice of our governmental representatives, but its effect has been extreme – we have exactly one African-American Senator (just barely) out of 100, and Barack Obama will be the first black president in the 220 years since Washington's first election. Obama will be living in a White House that was built by slaves. It would be great to be post-racial, but we're not; this is still a big deal. I can only hope that it will mark the beginning of an era when what really matters is the long-term success and happiness for us all, and that our grandchildren and great-grandchildren will be a lot smarter than we have been. Despite the enormous problems listed at the start of this piece, there is a collective vision of peace and prosperity that we not only can strive for, we must.

Speaking of vision, this issue's featured research area is our new effort in Computer Vision. Professor Trevor Darrell is a newcomer to ICSI, and you'll learn more about him and his work in these pages. This area is a great complement to our existing work in other perceptual modalities (speech and text). I hope you enjoy reading about it.

news briefs

ICSI Director Nelson Morgan and Chairman of the Board Prof. Shankar Sastry at the 20th Anniversary

In October, ICSI celebrated its **20th anniversary** with a day-long celebration including talks by **Prof. Richard Karp**, **Prof. Vern Paxson**, **Prof. Krste Asanovic**, and a keynote presentation by **Ben Gomes** of Google.

Prof. Wolfgang Wahlster and Dr. Manfred Dietrich with seven past and present DAAD postdocs at the 20th Anniversary

Recent research on the profitability of spam, conducted by ICSI **Networking Group** researchers along with a team from UCSD, was featured in a Washington Post blog article on November 6, 2008.

Professor **Richard M. Karp**, head of ICSI's Algorithms Group, is the 2008 winner of the Dickson Prize in Science. According to the prize website, "The Dickson Prize in Science is awarded annually to the person who has been judged by Carnegie Mellon University to have made the most progress in the scientific field in the U.S. for the year in question." The prize will be presented to Professor Karp on March 25, 2009.

Professor **Nelson Morgan**, ICSI's Director, has been named to the Advisory Council of the International Speech Communication Society (ISCA). ISCA is the primary organization devoted to speech communication science and technology.

Professor **Richard M. Karp** was profiled in an editorial at Investors.com, and also in The Berkeley Science Review, Issue 15. The BSR feature is written by ICSI Speech Group grad student Dan Gillick, a regular contributor to the publication.

ICSI has had the privilege of overseeing six doctoral theses in the period covered by this newsletter. The Speech group graduated Drs. **Kofi Agyeman Boakye** and **David Gelbart**. The AI group graduated Drs. **Nancy Chang**, **Joe Makin**, **Eva Mok**, and **Steve Sinha**. More information on their theses can be found in the Publications section, beginning on page 9.

The Vision Group recently welcomed two new arrivals - Torsten Darrell, son of **Trevor Darrell** and his wife Lisa, on December 18th, and Nika Tio, daughter of graduate student **Kate Saenko**, on October 25th. Congratulations to both families, and welcome Torsten and Nika.

Above: Torsten Darrell, Below: Nika Tio

ICSI's annual **BEARS Open House** was held on February 12th from 2:00-5:00 p.m. Professor **Trevor Darrell**, head of the new Vision Group, presented a talk on his latest computer vision work, **Luke Gottlieb** showed a video demo of recent speaker diarization work by members of the Speech Group, and the **FrameNet** team demonstrated their latest improvements, including some work on different languages. In addition, representatives from all research groups presented posters summarizing recent results from various current research projects. This Open House is held in conjunction with UC Berkeley's BEARS 2009 (Berkeley EECS Annual Research Symposium).

Prof. Trevor Darrell of the Vision Group gives the keynote speech at ICSI's BEARS Open House

Dilek Hakanni-Tür of ICSI's Speech Group and Guiseppe Riccardi of the University of Trento have won an IEEE Signal Processing Society (SPS) 2008 Best Paper Award for their paper, "Active Learning: Theory and Applications to Automatic Speech Recognition" which appeared in the July 2005 issue of IEEE Transactions on Audio, Speech, and Language Processing. They will be presented with the award at the ICASSP 2009 conference in Taipei, Taiwan, April 19-24, 2009.

Continued from Page 1

which ways for comparative facial recognition, e.g. face shape, distance between the eyes, cheekbone position, etc., and identify mathematically significant correlations between faces that look alike. Once these features are identified, the group will construct algorithms that will help identify the target child based on similarities and differences between the reference pictures.

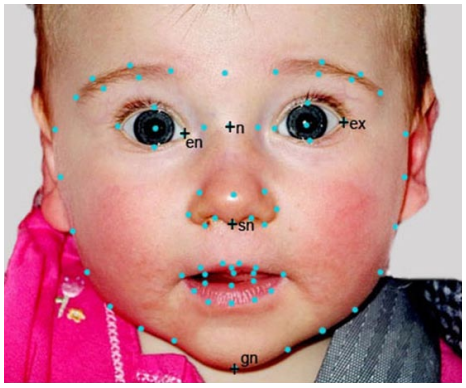

Nose-face height: n-sn/n-gn
Intercanthal index: en-en/ex-ex
An example of features analyzed for child facial recognition in the REUNITE project, work done by Ashley Eden. Photo taken and altered under the Creative Commons License from Flickr user 'ebruli.'

Hurricane Katrina exposed the frightening lack of emergency infrastructures. After the 2005 hurricane, it was 6 months until the last child was reunited with her family. In part, this is because there is no method for hospitals to document their unidentified children, and help parents search for them. Darrell's REUNITE project can dramatically augment rescue and relief efforts in large-scale disaster situations.

ROBOT VISION

Science fiction has filled our heads with the dream of human-like robots that can perform tasks ranging from mundane cleaning to world domination. Before functioning interactive robots could become a practical reality, researchers have a number of hurdles to overcome. Among these hurdles is the task of implementing robust robot vision, which members of the Vision group are developing with sponsorship from NSF, Toyota, and Menlo Park research start-up Willow Garage.

While it's easy for humans to identify everyday objects such as coffee cups, this poses a challenge for machines. Coffee cups can be made of a number of different materials, fall within a wide range of sizes, have a number of different shapes, come with or without various handle designs, and can be any color – including clear. With such a wide ranging set of variables for even a simple object, a robot has to be able to determine a lot of information through its visual sensors.

One thing working in the robot researcher's favor is the plethora of tagged images available online. With a simple Google image search for "coffee cup" returning nearly five million images, there is no shortage of training data. These data are not, however, of the most helpful quality. Current robot technology incorporates multiple sensors that utilize multi-spectral imaging to include qualities such as depth perception and infrared light information; information that is not present in the supply of internet images. This more complex multi-modal data is the most useful thing for training a robot to recognize an object, but there is not nearly as much information of this kind available. Darrell's group, along with researchers at Brown University and Willow Garage, is currently bridging the gap between the copious amount of web images and richer multi-modal data to allow robots to reap the benefits of both types of information and apply them to object recognition.

This brings us to IKEA; Darrell recently spent $300 and the better part of an afternoon there, purchasing one of every glass bowl and plate he could find. The group is building a test set of everyday objects to be catalogued using multi-spectral sensors, and developing a system which allows a robot to use high-level sensor information alongside the wealth of tagged web images. They are studying glassware for a number of reasons: the variability of individual items in a class mentioned above, so it makes for a good data set; robot vision currently has particular trouble with transparent objects, so it has the potential to work out another challenge for the field; and these are objects that household service

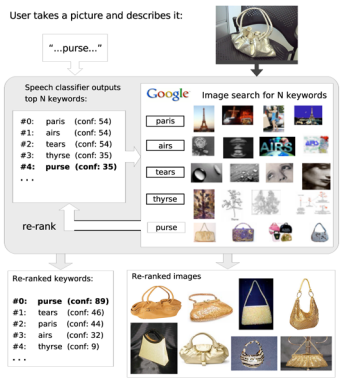

A representation of how a computer can learn visual sense classifiers without human supervision.

robots will commonly need to interact with, so it will be directly applicable. The group hopes that the training data from using the glassware will inform future efforts of how to reconcile multi-modal sensor input with the abundance of tagged web images. It won't be too long before there's a robot you can ask to "bring me my Peet's coffee cup from the dining room table", and it will be able to!

STRATEGIC MILITARY PLANNING

The Vision group has also been contracted to do some work for the US Military. The group has undertaken two projects that both involve utilizing real-time sensor data in strategic Military planning. One project, URGENT, generates three-dimensional maps on the fly; the other project, ULTRA-Vis, uses a technique called "augmented reality" to enhance human vision.

DARPA's URGENT program aims to take sensor feeds, such as those from satellites or helicopters, and extract three-dimensional information to create highly detailed maps in near-real-time. When the work is complete, sensors will be able to detect a wide array of elements and objects in an image, and extract the information for use in automated map generation. ICSI's Vision group is developing the algorithms that will allow this system to recognize objects such as stop signs, power lines, bus stops, and even garbage cans. The algorithm works with three-dimensional imagery to accurately identify objects that are occluded, or blocked from direct view. The overall system then could take this object information and combine it with terrain and building features to generate a map within half an hour, allowing for fast and highly-accurate planning of operations. The level of detail these maps can provide is not dissimilar from Google's wire-frame maps of downtown San Francisco that contain transit information and other points of interest.

ULTRA-Vis aims to improve soldier functionality in urban areas. Military personnel are faced with many challenges when operating in urban environments; there are often large non-combatant populations in locations where noise makes radio communication difficult, and buildings obscure line-of-sight, thereby making gesture communication difficult. The Vision group's portion of this project involves gesture recognition on wearable sensors; they are developing machine learning algorithms for automatic recognition of soldiers' hand gestures. Personnel will have wearable computers with see-through displays that will augment their vision by overlaying the information over what they are already seeing.

GROUP FUTURE

Even with the innovative work currently underway, there are still huge leaps to be made in computer vision. Besides the difficulty with transparent objects mentioned above, there are large conceptual hurdles to overcome, as well. One application of computer vision that is seeing rapid development is in mobile phones. Although currently limited by sensor quality and onboard processing power, mobile phone image recognition is seeing success with two-dimensional, high-contrast items like posters and CD covers as well as popular landmarks. Advances in phone-based image recognition will require efficient algorithms and/or low-latency off-board processing.

> "It is great to have Trevor join ICSI and UC Berkeley. His group has done some of the most interesting work on object recognition and on multimodal user interfaces, and this will make our effort in computer vision even stronger than before. We are already collaborating on several projects."
>
> -Prof. Jitendra Malik, UCB

While current object recognition algorithms require some kind of training or reference data, Darrell hopes to teach computers to understand the essence of what they're looking at. If a computer knows what properties constitute a chair, it will be able to make an intelligent – and more importantly, independent – determination that an object is or is not categorically a chair. The accuracy of this method would ideally parallel our own ability to recognize novel objects. When this kind of computational reasoning combines with affordable and highly portable sensors such as those on cell phones, there will be a wealth of information available about anything, just by pointing a camera at it.

notes from abroad: srini narayanan

Letter from the Institute for Advanced Study (IAS), Berlin

As a graduate student and then as a researcher at ICSI, I have had the pleasure of welcoming and working closely with many international researchers and have sometimes wondered what ICSI is like from the viewpoint of a visitor. Now I have a better perspective. For the first time in a decade, I have been away from Berkeley for over a month. As a resident fellow at the Institute for Advanced Study in Berlin, I am spending this semester in Berlin consolidating and writing up several years of work within the NTL project on neurally plausible computational models of cognitive phenomena.

I had not visited Berlin before, did not speak German except for a few rudimentary phrases, and was in all respects even less prepared for the transition than several of our visitors to ICSI. The administrative staff at the IAS managed to make the move completely smooth. They arranged for the residence permit (you need one for stays of over three months), a great apartment, office space, computers, supplies and even gave us invaluable tips on moving around, getting basic things done (like opening a bank account, groceries). All this could have been quite onerous, and the ease with which this happened made me appreciate the great service the corresponding ICSI staff provides to our numerous international vistors every year.

The Institute for Advanced Study (IAS) in Berlin (aka Wissenschaftskolleg zu Berlin or just Wiko) offers the selected researchers (around 40/year who stay anywhere between 6-10 months), the opportunity to concentrate on their research projects and to absorb ideas and inspiration from other disciplines. In selecting its members, the international advisory board of the IAS places no restrictions on country of origin, discipline, or academic position. The Fellows' only obligations are residence at the Wissenschaftskolleg and the requirement to meet once a day for a meal and each Tuesday for the weekly Colloquium. At each Colloquium, one Fellow presents his work to the others and all Fellows, regardless of their disciplinary background, consider and discuss the presented topic.

The intellectually heterogeneous atmosphere has been very instructive and has resulted in critical examination of past work as well as new problems to contemplate. In consolidating research within NTL, I have been reading a number of articles on the experimental techniques (such as fMRI imaging) used in our own work within NTL and in the wider context of cognitive neuroscience. This has resulted in an idea for a new experimental approach designed to move from the one-off studies that currently constitute the state of the field to a more predictive model that compares predicted fMRI results to actual images in test data. One derailment to my carefully laid plans involves fellow Fellow Robert Trivers, who has succeeded in getting me thinking about self deception and its connection to the evolution of deception, framing, and metaphor.

Berlin is an interesting and culturally rich city with great museums, nightlife, and real seasons (read: cold and grey winters). My experience and appreciation of western classical opera and musical concerts have benefited immensely even in the few months I have been here. As elsewhere in Europe, the ease with which one can use public transportation to explore this fairly large city and its surroundings is refreshing. The service economy, innovative green technologies, and the cosmopolitan nature of Berlin bears some similarities to the Bay area. Having said that, the difference in the quality of Indian food is substantial. So I have always found reasons to come back to Berkeley/ICSI every couple of months.

The pace here is much gentler than at ICSI, the conversations longer with emphasis placed on developing collaborations and entertaining perspectives on your work from scholars in diverse disciplines (history, biology, philosophy, law, art, music) who are often drinking wine during this interaction.[1] Clearly, one can tap the unparalleled resources at ICSI and on campus to achieve this degree of interdisciplinarity (and wine drinking is at least as refined an art in the bay area). But it does feel nice to have the legwork done and all the participants assembled in advance. On the other hand, I do miss the students to interact with, the research group meetings, and just the intellectual stimulation of several faculty, researchers, and students working toward a common objective. It would be nice to create a mini IAS within ICSI where we could have semester long international collaborations focused on topics of fundamental interest. This way we can have the best of both worlds. And in addition, good weather and better Indian food.

1 I have been noticing a corresponding increase in my coffee drinking, perhaps as a consequence.

featured alum: jeff bilmes

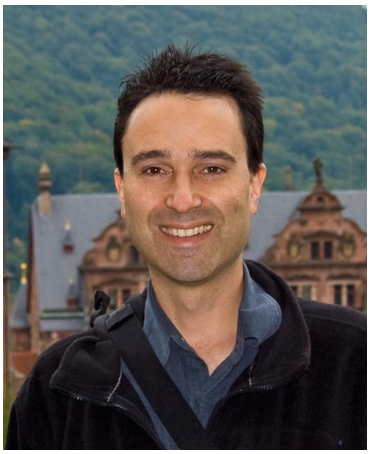

Professor Jeff Bilmes, of the University of Washington, is our featured alumnus for this issue. Bilmes was one of ICSI's very first employees, joining our ranks while a computer science undergraduate at U.C. Berkeley in the late 80's. Bilmes started out working as a programmer on spatial data-base access methods for Oliver Günther, but as staff members began to realize his unusually broad set of talents and interests, he was offered a research position in the Realization Group (now the Speech Group). He worked with them for a year before heading to MIT for his Master's degree and then returning to ICSI and UCB to complete his Ph.D. in Computer Science.

While at ICSI, Bilmes worked in a wide range of disciplines. His work with Prof. Nelson Morgan introduced entirely new methods to existing speech recognition algorithms. His early pioneering work on optimized numerical library auto-tuning, along with Krste Asanovic (another ICSI-to-MIT and back-to-ICSI researcher), generated the PHiPAC project, a method that could automatically generate code that would specifically tailor numerical programs to run at near-peak performance. This work spawned the entire field of auto-tuning in high-performance computing.

In 1999, Bilmes accepted a faculty position at the University of Washington, Seattle, where he received an NSF Career award, a CRA digital government fellowship, and a NAE Gilbreth Lectureship award. He has continued in his tradition of making breakthroughs in a number of different fields. At UW, he developed the Graphical Models Toolkit (GMTK), a software system that researchers can use to express and compute with an unlimited number of statistical models over sequential data, such as speech, language, and biological signals. GMTK allows computation with more sophisticated structures than simple Hidden Markov Models. GMTK is currently the most widely used graphical-model toolkit for sequential data, and is in heavy use by laboratories all over the world. Bilmes has published in a diverse set of areas, including speech recognition, natural language processing, human-computer interaction, bio-informatics, social networks, computer vision, active learning, semi-supervised learning, and submodularity in machine learning.

One particular project of Bilmes's that has enjoyed considerable media attention is the Vocal Joystick. Named one of the 25 leading-edge IT research projects in 2008 by Network World, the Vocal Joystick's aim is to enable individuals with motor impairments to control on-screen devices (like mouse pointers, scroll bars, and general computer interfaces) and electro-mechanical devices as fluently as any non-impaired user might do. With the Vocal Joystick system, a user uses simple and easy-to-learn continuous non-verbal vocalizations to control a mouse cursor (or a robotic arm). For example, instead of saying things like "go, up, left, stop", the Vocal Joystick uses these continuous vocalizations to cause a sweep from "up" to "left" at a manner and rate as fast or slow as the user desires. Applications of the Vocal Joystick range from basic mouse control, to sophisticated cad-design and/or photo touch-up applications, to art, and even to video games (as entertainment too should be made accessible to the motor impaired). Studies have shown that the Vocal Joystick greatly outperforms other voice-control methods in terms of output, experience, efficiency, and user satisfaction.

Bilmes is currently on a year-long sabbatical at the Max Planck Institute for Biological Cybernetics in Tübingen, Germany. He's in Prof. Dr. Bernard Schölkopf's group working on various writing and other projects related to machine learning, robotics, and sequential graphical models.

> "Science, engineering, programming, you name it; he was always at what the baseball guys would call '110%'."
>
> -Prof. Nelson Morgan

Of his time at ICSI, Director Morgan says Bilmes "was in permanent overdrive. Science, engineering programming, music, you name it; he was always at what the baseball guys would call '110%'." Prof. Bilmes is a great example of the kind of researcher that ICSI tries to foster. He is still collaborating internationally, and is just as comfortable crossing disciplines as walking down the hall.

visiting scholars


Dorothea Emig

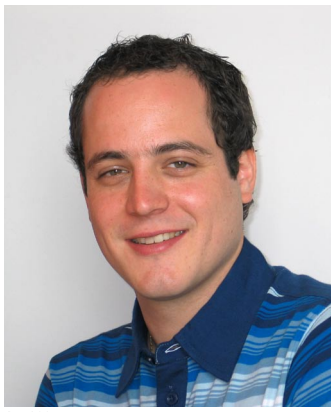
David Imseng

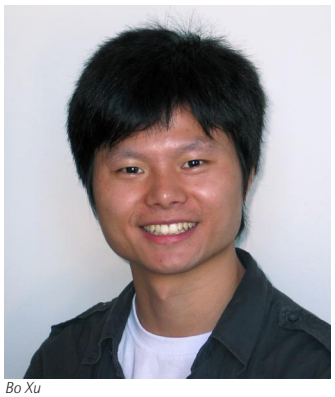
Bo Xu

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

From China
Bin Dai (Networking)
Bo Xu (Networking)

From European Union (AMIDA)
Carl Henrik Ek (Vision)
Korbinian Riedhammer (Speech)
Shasha Xie (Speech)

From Finland
Yoshia Hirase (Industrial)
Jyri Kivinen (Algorithms/Campus Affiliation)
Kimmo Kuusilinna (Industrial)
Tiina Lindh-Knuutila (AI)
Tommi Lampikoski (Campus Affiliation)
Annukka Näyhä (Campus Affiliation)
Boris Nechaev (Networking)

From Switzerland (IM2)
Nikhil Garg (Speech)
David Imseng (Speech)

From Germany
Jan Baumbach (Algorithms)
Gerald Friedland (Speech)
Tobias Friedrich (Algorithms)
Martin Gairing (Algorithms)
Oliver Guenther (Industrial)
Birte Lönneker-Rodman (AI-FrameNet)
Ulrich Rueckert (Algorithms)
Felix Salfner (Industrial)
Thomas Sauerwald (Algorithms)
Guido Schryen (Networking)
Holger Ziekow (Other)

From Spain
Oscar Ferrandez (AI)
Eduardo Lopez (Speech)
Carlos Subirats (AI-FrameNet)
Enrique Torres (Architecture)

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

AI
Alberto Amengual
Terry Regier

Algorithms
Dorothea Emig
Ron Shamir
Tobias Wittkop

Extreme Architecture
Chris Batten
Rose Liu

FrameNet
Kyoko Ohara
Hiroaki Sato

Networking
Po-Ching Lin
Jamon Liu
Chris Kanich
Gregor Maier

Other
Oliver Günther
Agnés Voisard

publications listing

L. Aziz-Zadeh, C. Fiebach, S. Narayanan, J. Feldman, E. Dodge, and R. B. Ivry. Modulation of the FFA and PPA by Language Related to Faces and Places. Social Neuroscience, Vol. 3, Issues 3-4, pp. 229-238, September 2008.

S. Baswana, S. Biswas, DB. Doerr, T. Friedrich, P. Kurur, and F. Neumann. Computing Single Source Shortest Paths Using Single-Objective Fitness Functions. Proceedings of the ACM International Workshop on Foundations of Genetic Algorithms (FOGA 2009), Orlando, Florida, January 2009.

C. Batten, H. Aoki, and K. Asanovic. The Case for Malleable Stream Architectures. Proceedings of Workshop on Streaming Systems at the 41st International Symposium on Microarchitecture (MICRO-41), Lake Como, Italy, November 2008.

J. Baumbach, S. Rahmann, and A. Tauch. Reliable Transfer of Transcriptional Gene Regulatory Networks Between Taxonomically Related Organisms. BMC Systems Biology, Vol. 3, Issue 8, January 2009.

J. Baumbach, S. Rahmann, and A. Tauch. Towards the Integrated Analysis, Visualization, and Reconstruction of Microbial Gene Regulatory Networks. Briefings in Bioinformatics, Vol. 10, Issue 1, January 2009.

P. Bensel, O. Günther, C. Tribowski, and S. Vogeler. Cost-Benefit Sharing in Cross-Company RFID Applications. Proceedings of International Conference on Information Systems (ICIS), Paris, France, paper 129, December 2008.

K. A. Boakye. Audio Segmentation for Meetings Speech Processing. UC Berkeley dissertation, December 2008.

K. Bringmann and T. Friedrich. Approximating the Volume of Unions and Intersections of High-Dimensional Geometric Objects. Proceedings of International Symposium on Algorithms and Computation (ISAAC), Gold Coast, Australia, pp. 436-447, December 2008.

K. Bringmann and T. Friedrich. Don't Be Greedy When Calculating Hypervolume Contributions. Proceedings of the ACM International Workshop on Foundations of Genetic Algorithms (FOGA 2009), Orlando, Florida, January 2009.

N. Chang. Constructing Grammar: A Computational Model of the Emergence of Early Constructions. UC Berkeley dissertation, February 2009.

B-G. Chun, P. Maniatis, S. Shenker, and J. Kubiatowicz. Tiered Fault Tolerance for Long-Term Integrity. Proceedings of USENIX Conference on File and Storage Technologies (FAST), San Francisco, California, February 2009.

J. DeNero, A. Bouchard-Côté, and D. Klein. Sampling Alignment Structure Under a Bayesian Translation Model. Proceedings of Conference on Empirical Methods in Natural Language Processing (EMNLP), Waikiki, Honolulu, Hawaii, pp. 314-323, October 2008.

B. Doerr, T. Friedrich, M. Künneman, and T. Sauerwald. Quasirandom Rumor Spreading: An Experimental Analysis. Proceedings of the Workshop on Algorithm Engineering and Experiments (ALENEX09), New York, New York, pp. 145-153, January 2009.

M. Ellsworth, R. Lee-Goldman, and R. Rhodes. Two Paradoxes of English Determination: The Construction of Complex Number Expressions. Presented at Western Conference on Linguistics (WECOL 2008), Davis, California, November 2008.

A. Ermolinskiy, D. Moon, B-G. Chun, and S. Shenker. Minuet: Rethinking Concurrency Control in Storage Area Networks. Proceedings of USENIX Conference on File and Storage Technologies (FAST), San Francisco, California, February 2009.

A. Ermolinskiy and S. Shenker. Reducing Transient Disconnectivity Using Anomaly-Cognizant Forwarding. Proceedings of the 7th ACM Workshop on Hot Topics in Networks (HotNets-VII), Calgary, Canada, pp. 91-96, October 2008.

B. Fabian, O. Günther, and S. Spiekermann. ONS Security. RFID Handbook: Applications, Technology, Security, and Privacy, S. Ahson and M. Ilyas, eds., CRC Press, 2008.

B. Favre, D. Hakkani-Tur, S. Petrov, and D. Klein. Efficient Sentence Segmentation Using Syntatic Features. Proceedings of IEEE Workshop on Spoken Language Technologies (SLT2008), Goa, India, pp. 77-80, December 2008.

J. Feldman and S. Narayanan. Simulation and Cognition in Language Understanding. Proceedings of Exciting Biologies 2008, Chantilly, France, October 2008.

C. Fillmore. Review Article: A Valency Dictionary of English. International Journal of Lexicography, October 2008.

A. Franklin, G. V. Drivonikou, L. Bevis, P. Kay, and T. Regier. Categorical Perception of Color is Lateralized to the Right Hemisphere in Infants, but to the Left Hemisphere in Adults. Proceedings of the National Academy of Sciences, Vol. 105, Issue 9, pp. 3221-3225, March 2008.

A. Franklin, G. V. Drivonikou, A. Clifford, P. Kay, T. Regier, and I. R. L. Davies. Lateralization of Categorical Perception of Color Changes with Color Term Acquisition. Proceedings of the National Academy of Sciences, Vol. 105, No. 47, pp. 18221-18225, November 2008.

G. Friedland, H. Hung, and C. Yeo. Multi-Modal Speaker Diarization of Real-World Meeting Using Compressed-Domain Video Features. ICSI Technical Report TR-08-007, October 2008.

G. Friedland, W. Hürst, and L. Knipping. Multimedia Education—Can We Find Unity in Diversity? Proceedings of the 16th ACM International Conference on Multimedia, Vancouver, Canada, pp. 1115-1116, October 2008.

G. Friedland, L. Knipping, and W. Hürst. Automated Lecture Recording. Encyclopedia of Multimedia, B. Furht, ed., October 2008.

G. Friedland and C. Martell, eds. Best Papers from the Second IEEE International Conference on Semantic Computing (IJSC). International Journal on Semantic Computing (IJSC), Vol. 2, Issue 3, September 2008.

G. Friedland and O. Vinyals. Live Speaker Identification in Conversations. Proceedings of the 16th ACM International Conference on Multimedia, Vancouver, Canada, pp. 1017-1018, October 2008.

Continued on Page 10

publications

Continued

T. Friedrich and N. Hebbinghaus. Average Update Times for Fully-Dynamic All-Pairs Shortest Paths. Proceedings of International Symposium on Algorithms and Computation (ISAAC), Gold Coast, Australia, pp. 693-704, December 2008.

M. Gairing. Malicious Bayesian Congestion Games. Proceedings of the 6th Workshop on Approximation and Online Algorithms (WAOA'08), Universität Karlsruhe, Germany, pp. 119-132, September 2008.

M. Gairing, T. Lücking, M. Mavronicolas, B. Monien, and M. Rode. Nash Equilibria in Discrete Routing Games with Convex Latency Functions. Journal of Computer and System Sciences, Vol. 74, Issue 7, pp. 1199-1225, November 2008.

N. P. Garg, S. Favre, H. Salamin, D. Hakkani-Tür, and A. Vinciarelli. Role Recognition for Meeting Participants: An Approach Based on Lexical Information and Social Network Analysis. Proceedings of the 16th ACM International Conference on Multimedia, Vancouver, Canada, pp. 693-696, October 2008.

N. Garg and I. Weber. Personalized, Interactive Tag Recommendation for Flickr. Proceedings of the 2nd ACM International Conference on Recommender Systems, Lausanne, Switzerland, pp. 67-74, October 2008.

D. Gelbart. Ensemble Feature Selection for Multi-Stream Automatic Speech Recognition. UC Berkeley dissertation, December 2008.

A. Gilbert, T. Regier, P. Kay, and R. Ivry. Support for Lateralization of the Whorf Effect Beyond the Realm of Color Discrimination. Brain and Language, Vol. 105, Issue 2, pp. 91-98, May 2008.

D. Gillick, B. Favre, and D. Hakkani-Tur. The ICSI Summarization System at TAC 2008. Proceedings of Text Analysis Conference (TAC), Gaithersburg, Maryland, November 2008.

P. B. Godfrey, S. Shenker, and I. Stoica. Pathlet Routing. Proceedings of the 7th ACM Workshop on Hot Topics in Networks (HotNets-VII), Calgary, Canada, pp. 97-102, October 2008.

C. Goebel and O. Günther. The Value of Visibility and RFID in Transshipment Operations—A Simulation Study. Proceedings of the 9th Internationale Tagung Wirtschaftsinformatik (WI 2009) (Conference on Business Informatics), Vienna, Austria, February 2009.

C. Goebel, H. Krasnova, T. Hildebrand, O. Günther, and F. Bidault. Enterprise Software Adoption and Perceived Business Performance—An Empirical Investigation in Germany and the U.S. Proceedings of International Conference on Information Systems (ICIS), Paris, France, paper 57, December 2008.

C. Goebel, H. Krasnova, H. Syllwasschy, and O. Günther. Vertical Integration and Information Sharing—An Empirical Investigation in the German Apparel Industry. Proceedings of the 9th Internationale Tagung Wirtschaftsinformatik (WI 2009) (Conference on Business Informatics), Vienna, Austria, February 2009.

C. Goebel, C. Tribowski, and O. Günther. EPCIS-Based Supply Chain Event Management—A Quantitative Comparison of Candidate System Architectures. Proceedings of International Conference on Complex, Intelligent, and Software Intensive Systems (CISIS-2009), Fukuoka, Japan, March 2009.

P. V. Gopalacharyulu, V. R. Velagapudi, E. Lindfors, E. Halperin, and M. Oresic. Dynamic Network Topology Changes in Functional Modules Predict Responses to Oxidative Stress in Yeast. Molecular BioSystems, Vol. 5, Issue 3, pp. 276-287, 2009.

S. Gürses, R. Rizk, and O. Günther. Privacy Design in Online Social Networks: Learning from Privacy Breaches and Community Feedback. Proceedings of International Conference on Information Systems (ICIS), Paris, France, paper 90, December 2008.

H. Hung and G. Friedland. Towards Audio-Visual On-Line Diarization of Participants in Group Meetings. Proceedings of European Conference on Computer Vision (ECCV), Marseille, France, October 2008.

C. G. Jones, R. Liu, L. Meyerovich, K. Asanovic, and R. Bodik. Parallelizing the Web Browser. Proceedings of the First USENIX Workshop on Hot Topics in Parallelism (HotPar'09), Berkeley, California, March 2009.

C. Kanich, C. Kreibich, K. Levchenko, B. Enright, G. Voelker, V. Paxson, and S. Savage. Spamalytics: An Empirical Analysis of Spam Marketing Conversion. Proceedings of the 15th ACM Conference on Computer and Communications Security (ACM CCS), Alexandria, Virginia, pp. 3-14, October 2008.

P. Kay and L. A. Michaelis. Constructional Meaning and Compositionality. Semantics: An International Handbook of Natural Language Meaning, HSK Handbooks of Linguistics and Communication Science Series: 23: Semantics and Computer Science, C. Maienborn, K. von Heusinger, and P. Portner, eds., Mouton de Gruyter, 2008.

P. Kay, T. Regier, A. L. Gilbert, and R. B. Ivry. Lateralized Whorf: Language Influences Perceptual Decision in the Right Visual Field. Language, Evolution, and the Brain, J. W. Minett and W. S-Y. Wang, eds., City University of Hong Kong Press, 2009, in press.

G. Kimmel, R. M. Karp, M. A. Jordan, and E. Halperin. Association Mapping and Significance Estimation via the Coalescent. American Journal of Human Genetics, Vol. 83, Issue 6, pp. 675-683, December 2008.

R. Lee-Goldman. Infinitives in Comparatives: Canonical Syntax Meets Quirky Semantics. Presented at Annual Meeting of the Linguistic Society of America, San Francisco, California, January 2009.

R. Lee-Goldman, R. Rhodes, and M. Ellsworth. Determination and Modification: Interaction and Interpretation. Presented at the 83rd Annual Meeting of the Linguistic Society of America, San Francisco, California, January 2009.

Z. Li, A. Goyal, Y. Chen, and V. Paxson. Automating Analysis of Large-Scale Botnet Probing Events. Proceedings of ACM Symposium on Information, Computer, and Communication Security (ASIACCS 2009), Sydney, Australia, March 2009.

R. Liu, K. Klues, S. Bird, S. Hofmeyr, K. Asanovic, and J. Kubiatowicz. Tessellation: Space-Time Partitioning in a Manycore Client OS. Proceedings of the First USENIX Workshop on Hot Topics in Parallelism (HotPar'09), Berkeley, California, March 2009.

J. Makin. A Computational Model of Human Blood Clotting: Simulation, Analysis, Control, and Validation. UC Berkeley dissertation, December 2008.

A. Mandal, D. Vergyri, W. Wang, J. Zheng, A. Stolcke, G. Tür, D. Hakkani-Tür, and N. Fazil Ayan. Efficient Data Selection for Machine Translation. Proceedings of IEEE/ACL Workshop on Spoken Languge Technologies (SLT), Goa, India, pp. 261-264, December 2008.

E. Mok. Contextual Bootstrapping for Grammar Learning. UC Berkeley dissertation, January 2009.

D. Moon, M. Casado, T. Koponen, and S. Shenker. Rethinking Packet Forwarding Hardware. Proceedings of ACM Special Interest Group on Data Communications Workshop on Hot Topics in Networks (HotNets-VII), Calgary, Canada, pp. 1-6, October 2008.

P. Motlicek, S. Ganapathy, H. Hermansky, H. Garudadri, and M. Athineos. Perceptually Motivated Sub-Band Decomposition for FDLP Audio Coding. Proceedings of 11th International Conference on Text, Speech, and Dialogue (TSD 2008), Brno, Czech Republic, pp. 435-442, September 2008.

G. Myers, G. Tür, L. Voss, B. Bolles, S. Kajarekar, E. Shriberg, and D. Hakkani-Tür. Multimedia Information Extraction Roadmap. Proceedings of the AAAI Fall Symposium on Multimedia Information Extraction, Arlington, Virginia, November 2008.

S. Sinha. Answering Questions about Complex Events. UC Berkeley dissertation, December 2008.

H. Pan, B. Hindman, and K. Asanovic. Lithe: Enabling Efficient Composition of Parallel Libraries. Proceedings of the First USENIX Workshop on Hot Topics in Parallelism (HotPar'09), Berkeley, California, March 2009.

S. Pfennigschmidt and A. Voisard. A Situation-Centric Approach to Meteorological Services in the SITUMET Platform. Proceedings of 16th ACM SIGSPATIAL International Conference on Advances in Geographis Information Systems (ACM GIS), Irvine, California, Article No. 66, November 2008.

A. Raabe and A. Felke. A SYSTEMC Language Extension for High-Level Reconfiguration Modeling. Proceedings of Forum on Specification, Verification, and Design Lanaguages (FDL 08), Stuttgart, Germany, pp. 55-60, September 2008.

A. Raabe, J. Tietjen, and J. K. Anlauf. An Exact and Efficient Triangle Intersection Test Hardware. Proceedings of International Conference on Computer Graphics Theory and Applications (GRAPP '09), Lisbon, Portugal, February 2009.

M. Raya, M. H. Manshaei, M. Felegyhazi, and J.-P. Hubaux. Revocation Games in Ephemeral Networks. Proceedings of 15th ACM Conference on Computer and Communications Security (ACM CCS), Alexandria, Virginia, pp. 199-210, October 2008.

T. Regier and M. Ali Khalidi. The Arab Street: Tracking a Political Metaphor. Middle East Journal, Vol. 63, Issue 1, pp. 11-29, 2009.

K. Riedhammer, B. Favre, and D. Hakkani-Tur. A Keyphrase Based Approach to Interactive Meeting Summarization. Proceedings of IEEE Workshop on Spoken Language Technologies (SLT2008), Goa, India, pp. 153-156, December 2008.

K. Saenko and T. Darrell. Unsupervised Learning of Visual Sense Models for Polysemous Words. Proceedings of the Twenty-Second Annual Conference on Neural Information Processing Systems (NIPS), Vancouver, Canada, December 2008.

K. Saenko, K. Livescu, J. Glass, and T. Darrell. Multistream Articulatory Feature-Based Models for Visual Speech Recognition. IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI), December 2008.

G. Schryen. An Analysis of Internet Voting Security: The Case of Estonia. Proceedings of the Workshop on e-Business (WEB 2008), Paris, France, December 2008, to appear. Also in Lecture Notes on Business Information Processing, Springer Heidelberg, 2009.

G. Schryen and R. Kadura. Open Source vs. Closed Source Software: Towards Measuring Security. Proceedings of the Annual ACM Symposium on Applied Computing, Honolulu, Hawaii, March 2009.

E. Shriberg, M. Graciarena, H. Bratt, A. Kathol, S. Kajarekar, H. Jameel, C. Richey, and F. Goodman. Effects of Vocal Effort and Speakign Style on Text-Independent Speaker Verification. Proceedings of Interspeech, Brisbane, Australia, pp. 609-612, September 2008.

E. Shriberg and A. Stolcke. The Case for Automatic Higher-Level Features in Forensic Speaker Recognition. Proceedings of Interspeech, Brisbane, Australia, pp. 1509-1512, September 2008.

S. Sinha. Answering Questions About Complex Events. UC Berkeley dissertation, December 2008.

G. Tür, A. Stolcke, L. Voss, J. Dowding, B. Favre, R. Fernandez, M. Frampton, M. Frandsen, C. Frederickson, M. Graciarena, D. Hakkani-Tur, D. Klintzing, K. Leveque, S. Mason, J. Niekrasz, S. Peters, M. Purver, K. Riedhammer, E. Shriberg, J. Tien, D. Vergyri, and F. Yang. The CALO Meeting Speech Recognition and Understanding System. Proceedings of IEEE Workshop on Spoken Language Technologies (SLT2008), Goa, India, pp. 69-72, December 2008.

O. Vinyals and G. Friedland. A Hardware-Independent Fast Logarithm Approximation with Adjustable Accuracy. Proceedings of the 10th IEEE International Symposium on Multimedia, Berkeley, California, pp. 61-65, December 2008.

N. Weaver, R. Sommer, and V. Paxson. Detecting Forged TCP Reset Packets. Proceedings of the 16th Annual Network and Distributed System Security Symposium (NDSS 2009), San Diego, California, February 2009.

R. Weiss, M. Mandel, and D. Ellis. Source Separation Based on Binaural Cues and Source Model Constraints. Proceedings of Interspeech, Brisbane, Australia, pp. 419-422, September 2008.

S. Yaman and C.-H. Lee. A Comparison of Single- and Multi-Objective Programming Approaches to Problems with Multiple Design Objectives. Journal of Signal Processing Systems, MLSP special issue, November 2008.

T. Yeh, J. Lee, and T. Darrell. Photo-Based Question Answering. Proceedings of 16th ACM International Conference on Multimedia 2008 (ACM Multimedia), Vancouver, Canada, pp. 389-398, October 2008.

INTERNATIONAL
COMPUTER SCIENCE
INSTITUTE

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

Editor-in-Chief
Maria Eugenia Quintana

Senior Editor
Nelson Morgan

Design
envision ID

Layout
Jacob Wolkenhauer

Copywriter
Jacob Wolkenhauer

INTERNATIONAL
COMPUTER
SCIENCE
INSTITUTE

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