The Computer Science Monitor

Tucson, AZ

VOL.III No.1 *SPRING 2011*

Head's Report

By SAUMYA DEBRAY

This newsletter celebrates the many activities in our department over last summer and fall. We start with an interview with Professor Greg Andrews, who recently became emeritus (and who just received a prestigious award from the University of Washington). Articles follow from a doctoral, MS, and undergraduate alumni: Gideon Myers, Alvin Gendrano, and Noah Snavely, respectively, the last interviewed by his brother, a current undergraduate in our department! Included are reports on department research (by Assistant Professor Chris Gniady and Research Faculty Ian Fasel) and a report on one of the departmental focal area (that of ergalics, by Professor Rick Snodgrass). We also happily feature our graduates for the last academic year, including abstracts from our doctoral graduates (Igor Crk, Scott Morris, and Joseph Schlecht), as well as a list of the recipients honored at our recent Department Awards Ceremony, a staff report on Academic Services, and last but not least, our thanks to our industrial partners and supporters.

This newsletter is published twice a year by the Publicity & Awards Committee
Computer Science Department, The University of Arizona
P.O. Box 210077, Gould-Simpson Building, Tucson, AZ 85721-0077
http://www.cs.arizona.edu

An interview with Greg Andrews

By PETER DOWNEY

Q: Emeritus! How does it feel so far?

Great!

Q: You got your BS in Math from Stanford in 69. When did you feel that you had become a computer scientist? There were not that many around in those days.

I took my first computing class sophomore year. When I arrived I had an idea of a major—math—but by sophomore year I was questioning my sanity and I got introduced to computer science. I just thought this was fascinating. But [Math] was the simplest major to finish up, and would let me take a variety of computer science classes. *Q: You went to the University of Washington and got your Ph.D. in 74. What were your original areas of interest, and how did they change throughout that period?*

I had a smattering. I did programming. I had numerical analysis and architecture and symbolic computing (Lisp) — and I liked programming, I liked problem solving so I think that drew me toward the system side of things. But practically nothing had a book. The field was being born. Knuth "Vol II" was a graduate textbook. I mean, what was really neat—was that we were at the dawn of the field, and so the books were being written. I took a compilers course with David Gries—he wrote the first compilers book and it was in draft form.

Q: Who were your role models?

The fascination of the field was the original one. But the most influential teacher I had was David Gries because he was a fantastic teacher and I got to talk to him quite a bit. David was certainly influential—and proved to be really influential later. In graduate school the person I had two classes from was Jean-Loup Baer—neither class was architecture: it was formal languages and data structures. He was the first hire at the University of Washington with a degree in Computer Science.

Q: I recall that my own dissertation—which must have been around the same time as yours—was written using 3 different IBM Selectric type balls. Output devices were non-existent. What was your format? What existed?

Nothing. It was typed. Alan Shaw was my adviser, I had a research assistantship with him and he paid for typing. The [final draft] was all handwritten.

Q: Your first job as a professor was Cornell. Who was there? What were the systems like when you arrived?

[Juris] Hartmanis, [John] Hopcroft, [Bob] Constable, David Gries. [Computing] was an IBM mainframe in the basement so OS 360 was the operating system. I first programmed on an IBM 7094 and then an IBM 360. But—punch cards!—all the way through—even at Cornell. I never typed a paper—I never had a terminal on my desk until I came to Arizona.

Q: I don't think you had ever heard of the University of Arizona before you came to visit in 1979. What attracted you to the University and the area?

The Department was small and collegial. The area was attractive. But you remember how you guys completely orchestrated things, so that [my wife] Mary saw the greenest parts of the city—we were ready to get out of the snow and the sleet. And of course [David] Dobkin and Ravi [Sethi]—and we had [Tim] Budd and [Chris] Fraser and [Dave] Hanson and Gary Levin.

Q: You stepped up as Department Head in 1986. What do you recall as the most significant events in the life of the Department during that 80s era?

The first NSF Infrastructure Grant had really changed the Department and made so many more things possible. And that was awarded in 84. And so 89 was extremely significant for two events: we did the [grant] renewal and had the infrastructure grant awarded in 89 and the undergraduate program started in the fall of 89. So those were the most significant events. The Department had stabilized, we had a solid core and we had the prospect of growth and then we hit the first of the cyclic recessions that we have hit—the one in 91—and that put the brakes on. We got the first 3 positions, staff, the Lab, and a huge increase in the operations budget. And then we started getting hit with budget cuts and freezes.

Q: In the early 00s, you served as an NSF Division Director, first in Experimental and Integrative Activities, and then Computer and Network Systems. What was the transition like when you arrived in Washington, D.C.?

Incredible learning curve. I spent the first few months every night preparing for the next day's meetings. I was learning forced methods because Integrative Activities involved a whole potpourri of things so I was learning a tremendous amount. And we were kicking around ideas for organizing the CISE (Computer Information Science and Engineering) Directorate, so that was a significant effort. It was an intellectual reorganization but then also a managerial or administrative reorganization. It gave us much more budgetary flexibility because we could target—if we got proposals in an area we could respond to where the interesting research problems were with much more flexibility.

Q: Just recently, University of Washington's Computer Science and Engineering department announced that you are one of two people being awarded the Outstanding Alumnus Award—a brand new award. My congratulations! Does UW keep close tabs on its alumni?

Yes they have had a newsletter and kept tabs on alumni for many years.

Q: You told me you spend about a day a week in the Department working on iPlant.

And advising—I still have two Ph.D. students. The nature of the grant is that it is a $50M award from the NSF from the Biology Directorate to build a computing infrastructure to support the solution of grand challenge problems in the biosciences. In the first year of the award through series of workshops with the bioscience community [we] identified two problems: A plant tree of life (there are 5000 green plant species—they would like to know how they are all related evolutionarily). The other one is called "from genotype to phenotype." The phenotype of a plant is its phenomena—what you see—its flowering, productivity in food production—what you observe. And the question there is: how does the genetic structure predict and affect the phenomena that result?

Anyway, I have two roles. My main role is as a member of the leadership advisory team. The second role is pushing the envelope on what one does with the computing infrastructure and that led to a project to find out: "How do you reproduce software experiments?" But have you ever tried to run a program that is five years old? You can't compile it, the language has changed,—and libraries have changed. This project is to build a set of tools that enable one to very simply, without changing the programs, capture that program and be able to run the program later—to be able to capture the experimental apparatus so that using the same apparatus one can run new experiments. *Q: The final question is: of all your experiences here over the years, what was the best?*

What's most important are the students—seeing students that you have mentored develop and go on to careers of their own. And since we started the undergraduate program we have such good undergraduate students there have been several undergraduate students who have done phenomenally well. The first undergrad I employed as a research assistant was Dawson Engeler who went on to a Ph.D. at MIT and is now a tenured professor at Stanford. This was a guy who didn't even have a 3.0 when he graduated! But I could see he loved research and was extremely good at it, and MIT took a chance!

Alumni Reports

Doctoral Alumnus: Gideon Myles

(Gideon is now an attorney at Novak Druce + Quigg LLP in San Francisco, CA.)

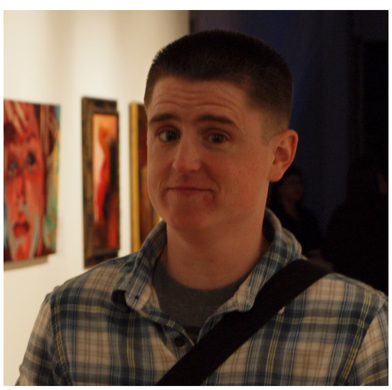

Since completing my PhD at the University of Arizona in May 2006 my career has taken a rather interesting, and definitely unexpected, path. Well, actually my career journey began even before I graduated. In the summer of 2004 I accepted a summer internship with the Content Protection group at the IBM Almaden Research Center. One of the areas that the Content Protection group was exploring was Software Protection, more specifically code obfuscation and tamper resistant code, which was a perfect fit with my PhD research. My original intention in taking the internship was to gain industry experience because I thought it would make me a better professor. Oh yeah, maybe I should mention that when I applied for the PhD program my career goal was to teach at a liberal arts college, preferably on the East Coast.

During my internship at IBM I had some great experiences, the highlight of which was inventing a new software watermarking algorithm that IBM decided to try to patent. About half way through my internship, my manager came to me to tell me that they were very interested in me staying on full time after my internship was over. This was a difficult decision to make for a number of reasons. First, I knew that I still had at least another year before I would be finished with my PhD and I had heard so many stories about people who had taken jobs prior to completing their dissertation. Second, I was worried that it might be taking me away from the career path I had laid out for myself. Finally, it meant I would have to move to California, which definitely was not on the East Coast. I ended up taking the job because I thought it would be a great opportunity to get more research and industry experience, which I thought would make me even more marketable once I was ready to apply for an academic position. Once, I completed my PhD I did decide to apply for a number of academic positions, as well as for a permanent position at IBM Research. Ultimately, I was offered two academic positions and a position as a Research Staff Member at Almaden. I decided to take the position at Almaden.

Working at IBM Research was a great experience because it was kind of a middle ground between Academia and Industry: I would get to do the research that I loved and get to have a real world impact. When I worked at IBM and people asked me what I did I would always tell them that I got paid to make stuff up. In reality that is what I did all day. While I probably didn't have quite as much freedom as academia, as long as I could make a business justification for it, I could work on it. At IBM I got paid to do research, write papers, go to conferences all over the world, and make presentations to many IBM customers. I had an opportunity to try and convince people that software protection was an extremely valuable business asset.

Even though I loved getting paid to invent new software protection techniques, I quickly grew frustrated with how slowly a company of 300,000 people moved. I was working in a very small group so it was very hard to get any traction. This meant that one of the main reasons that I took the position, to have a real world impact, probably wasn't ever going to happen. So that is when I decided to pursue an interest I had developed early in my dissertation research: the law. I applied to Santa Clara University School of Law and started taking night classes in Fall 2007. I thought that with my technical background and living in Silicon Valley, I had an opportunity to do some very interesting work in Intellectual Property Law.

Not long after I started law school, I was presented with what I thought was a dream job. Out of the blue I was contacted by a recruiter at Apple Inc. and after a few interviews with the Digital Rights Management Group I was offered a position. Now I was really faced with a dilemma. On the one hand, I had been an Apple fan since I was a very little kid. But, I also knew that I would definitely be moving off of the academic path and I didn't know if I could handle working for such a demanding company while going to school at night. Ultimately, I decided the opportunity was too great to pass up, so I was just going to have to make it work.

Taking the job at Apple was probably one of the best decisions I have made. Being inside of Apple is almost indescribable. You work extremely hard and sometimes put in very long days, but in the end I found it to be worth it because I was very proud of the products that we produced. I have to say there wasn't a single day that I wasn't happy to get up and go to work.

Over the course of 5.5 years I have worked for two major computer companies and they couldn't be more different. I wouldn't say one is better than the other, but I would say that one turned out to be a better fit for my personality. IBM was a very large, established, and traditional company where I could be very open about what I was working on, but where my work seemed to move very slowly and have little real world impact. On the other hand, Apple was a large, quickly moving, extremely secretive company, but my work was key to an extremely profitable part of Apple's business: iTunes.

After almost 2.5 wonderful years at Apple, I was again presented with a great opportunity, but this time it meant a complete career change. At the end of April I made the move to work at Novak Druce + Quigg LLP, a San Francisco law firm that specializes in Intellectual Property Protection, where I am working on patent prosecution. This is the portion of the law that takes an invention from patent application through patent issue and you are required to have a technical background. Until I pass the Patent Bar I will be working as a Technical Advisor, and then I will move up the ranks to Patent Agent. Once I complete law school in May 2011 and pass the California Bar I will be a full fledged Patent Attorney.

I think that what my rather short career illustrates is that there are many opportunities available to people with a Computer Science background. These opportunities may not be the ones contemplated when starting school and they might not even be traditional computer science jobs, but they can be very rewarding. Sometimes you have to be flexible and try things out as they come along and sometimes you have to take risks and start forging a new path.

M.Sc. Alumnus: Alvin Gendrano

By BONGKI MOON

(ALVIN IS NOW DIRECTOR, SQLSERVER BUSINESS MANAGEMENT AT MICROSOFT IN REDMOND, WA.)

Q: Since you graduated, how and what have you been doing?

Since graduating from UA in 1998, I've held several positions at Microsoft and have learned a lot along the way. My first role involved supporting critical enterprise deployments of SQLServer 6.5 and I quickly became a database expert for DBAs in need of help. This job was eye-opening as it broadened my view of data systems from the strong theoretical foundation I received in school with a sense of the practical and that data had more intrinsic value than its stewards. Building on practical experiences, I took a job as a developer consultant helping large third-Party Independent Software Vendors (ISVs) develop windows-based solutions for Financial Services customers. Unsurprisingly, structured information was always at the center of the solution and my database development background quickly paid off. I was soon flying around the nation delivering architectural reviews, workshops, performance tuning labs, and strategy reviews. Seeing value from the onsite consultative services for large ISVs, I then ran a program to deliver the same type of developer services to over 15,000 ISVs in the Microsoft Partner program via tele-consultants we hired in China, India, Japan, Americas, Europe, Middle East, and Africa. This opportunity allowed me to experience global team management and P&L ownership while completing an MBA.

I then decided to pursue my childhood fantasy of working for the games industry and got a job managing a significant portion of the Xbox LIVE Digital Distribution business which saw 150%+ growth annually in the two years I managed it. Last month, I moved back to my database roots and started as the director of business management for the SQL Server product.

Q: What is your typical day at Microsoft like?

A typical day is very collegial, filled with hallway discussions involving some of the smartest and most passionate people in the world with a diverse range of perspectives. Though sometimes, the same passions do end up in lively debate and occasional hurt egos. What always amazes me is how small the company still is as evidenced by how I would run into folks from prior groups almost every day, most of them doing other things from last we met years ago.

I've been part of the most exciting Microsoft groups (Xbox LIVE and SQLServer) and the organizational energy behind how these products are changing the world of entertainment and business is very evident and quite infectious.

Q: Is there any lasting influence or lesson you got from UA?

My time at UA was the first time that I really had to live alone and fend for myself without a family support structure around me. Though scary at first, the experience left me with a new sense of independence and a confidence that I could take care of myself and others. It's a true life-lesson that I've continued to build on with every geographic move, the purchase of our first house, and providing for a growing family.

Thanks to Professor Snodgrass, UA also helped orient me towards the wonders of database software technology and a greater appreciation for data. This passion for databases and information has helped me succeed in all the positions I've taken on in the last 12 years and will continue to help me as businesses become even more information-centric in the future.

Q: What is UA or your life at UA remembered best for?

My life at UA is best remembered for what was for me a new frontier of friendships, places, and technology. I can still picture the nighttime desert treks under clear starry skies, the Saguaro canyon hike to a swimming hole with friends, camping at Mt Lemmon, the Arizona-Sonora Desert Museum, and a bevy of other wonderfully memorable places. I also made international friendships at the UA, many of whom I'm still in contact with today. Lastly, the professors and the curriculum opened my eyes to what would continue to be relevant technology for decades to come. Especially memorable was our published research paper on Parallel Temporal Aggregates which we presented at the ICDE in Sydney, Australia. It was a lot of work, but I'm sure glad we invested the time to complete the paper and hope to add-on to the research in the future.

Q: What would you like to talk about most with grad students here now?

When I'm asked for mentoring advice and/or career guidance, I would always ask one question: "Where do you see yourself in 5 - 10 years?". More often than not, I would get a blank quizzed look at the first meeting and I would ask the person to come back after figuring it out. The answer is important as it is unproductive to figure out what to do next career-wise without an end-goal in mind. Think about it, the series of steps needed to start your own technology business is different from a career in development or in academe. A good long-term goal allows you to prune many branches from an otherwise complex decision tree. After pruning potential paths, it's easier to strategize on the next steps by looking at it as a series of chess moves, with each move designed to directionally get you closer to your end-state.

When I left the Philippines years ago, I had every intention of coming back home to help reverse a phenomena known as "Brain Drain". This has been a key theme that I've used to guide my career choices. Knowing that Microsoft in the Philippines had 3 types of employee groups: sales, marketing, and services, I set out to choose career options that would give me the experiences needed at a subsidiary office. I got services experience with the Microsoft Consulting Services (MCS), got sales experience and P&L ownership from Xbox LIVE, and now engaging with SQLServer marketing in my current role to round out the experiences I need. Key point is that you need a solid career strategy now and it starts by defining your long-term goal.

Undergrad Alumnus: Noah Snavely

By WILLIAM SNAVELY

(NOAH IS NOW AN ASSISTANT PROFESSOR AT CORNELL UNIVERSITY IN ITHACA, NY.)

Q: How did you become interested in computer science?

I've always like computers and computer games, and in high school I was fascinated by the idea of artificial intelligence. I had an off-the-wall idea that I would write a program that would learn how to converse in English by participating in online chat rooms and learning from interactions with real people—it sounded like an easy idea at the time. I didn't really think of computer science as a field of study, though, and I came to UA interested in doing astronomy. But at the same time, I took a few computer science classes and started to realized how deep of a subject it was, then quickly became a CS major. What I really liked about CS was that it represented an intriguing combination of problems ranging from the theoretical—how can we represent and reason about computing in the abstract?—to the applied—how can we implement this low-level system call without getting a blue screen of death?

I didn't really appreciate this when I was starting college, but beyond being a fascinating discipline in its own right, computing is an essential part of our society, and in the past five decades computer scientists have had a dramatic impact on the world. Now is an especially exciting time to be in CS—as information of all forms is becoming more and more available, computer science is in a great position to provide the foundation for this revolution in information.

Q: What you've been working on since you left the U of A? What are your current research goals?

I'm working on modeling and visualizing the world in 3D from the massive amount of visual data available on the Internet—the millions of tourist photos on Flickr, for instance—through new computer vision and graphics algorithms. You can only go so far with photos on the Internet, though, since people tend to take the same kinds of photos. Everyone takes a photo of the Colosseum, for instance, but no one bothers to take photos inside some random building in Rome. I'm currently working on a project called PhotoCity (`http://photocitygame.com`), where people earn points and prizes for uploading new viewpoints, as part of an online game.

I'm also teaching an honors CS100 class called "Introduction to Computing using Matlab and Robotics," where we teach Matlab programming and basic computer science through having students program robots to see and understand their world. You might not think it at first, but it's amazing how much computer science you can teach through the simple problem of taking an image and thresholding it to find a red object—through this one problem alone we cover sorting, selection, runtime analysis, graphs, connected components, and convex hull.

Q: What knowledge, skills, tools have proved especially valuable to you?

I think the most important skill I learned from my CS classes is how to solve problems. What we do every day in CS is figure out how to take a complicated real-world problem that might at first seem hard to wrap your head around, then formulate it in a way where it starts to make sense, decompose it into more tractable pieces, figure out what tools are best for each piece, and in the end be able to answer questions like "does this algorithm actually work?" and "will it take two seconds or 100 million years to get the answer?" This is the kind of thinking that I try to instill in my CS100 students.

Q: What motivated you to stay in academia?

It was a tough decision—industry has its own set of fascinating problems. In the end, I decided I could have more of a long-term impact in academia, both by working on problems I thought were interesting, but also by teaching and mentoring new generations of computer science.

Departmental Research

New Computer Optimizations Save Time and Energy and Prolong Battery Life

By CHRIS GNIADY

The Environmental Protection Agency estimates that 2% of total US electricity consumption is from PCs, while a Congressional analysis reports another 1.5% is from data centers. These percentages will increase every year as our computing infrastructure continues to grow. It is common for single family homes to be equipped with several computers used for school, work, and entertainment. Additionally, we rely on portable computers for our daily activities, from which we expect high performance and long battery life. Subsequently, high energy consumption manifests itself directly in energy costs that homes and businesses have to pay, as well as shorter battery lives in portable units. High energy consumption also manifests itself in rising carbon emissions that are of increasing importance as we try to preserve our environment for future generations. Energy conservation in computer systems is critical, and many research teams have undertaken the task of creating greener computing; but usually at the cost of frustrating delays observed by computer users. Waiting several seconds for a document to open or a spreadsheet to be sorted is not an optimal utilization of user time and may be detrimental to overall productivity.

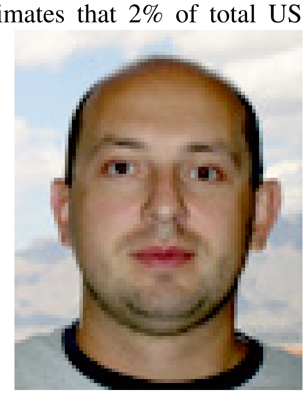
Chris Gniady

A team of researchers at the University of Arizona led by Chris Gniady have developed a novel system addition called Interaction Aware Prediction Mechanisms (IAPMs). IAPMs adapt the computer by understanding the interactions and expectations the user exhibits. Subsequently, the mechanisms are able to improve responsiveness, allowing users to accomplish tasks quickly, while saving energy, extending battery life, and lowering energy costs for desktop computers. For example, when a user attempts to open a file in a word processor, existing energy saving optimizations show a delay of one second. IAPMSs are able to notice several keystrokes or mouse clicks to detect user intent to open a file and deliver it quickly, even if the disk is in a deep power saving state.

The research team further extended the IAPMs to other system components such as the processor, memory, and network interface. All of these are responsible for performance and user experience, as well as the majority of energy consumption. For instance, the mechanisms detect the need to sort a large spreadsheet and provide the computing power and memory, allowing the processor to accomplish the task quickly without further delays. Alternatively, lower demand user needs for system performance such as typing, email, etc., can be detected by the IAPMs, allowing all devices to transition to a deep energy saving state. As such, users do not notice performance degradation when executing complex tasks, and gain longer battery life and lower monetary costs during the execution of simple tasks.

The critical observation of the research team at the University of Arizona is that the majority of functions place low demand on the system, rarely invoking complex tasks. Therefore the proposed IAPMs achieve excellent energy savings, while providing high performance when users need it. A combination of high performance and energy savings was a challenge for many researchers; the research team was able to solve the challenge without requiring users to alter their behavior or understanding of computer functionality. The proposed changes to computers will result in longer battery operation and lower electricity costs. Users will be able to use their computers just as before, and potentially apply them to new environments where long battery life is critical.

The Ergalics Focal Area

By RICHARD SNODGRASS

Ergalics is "the science of computational tools and of computation itself." (The phrase "natural science of computational tools as a perspective of computer science" is verbose and awkward. The term "science of computer science" is shorter but still awkward. Hence, we use the term "ergalics" for this new science. It derives from the Greek word ergaleion ($\epsilon\rho\gamma\alpha\lambda\epsilon\iota\omega\nu$), translated as "tool" or "instrument.")

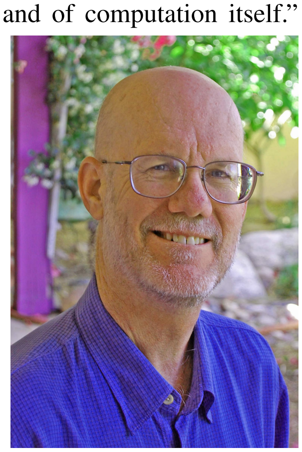
Rick Snodgrass

The articulation of scientific theories and the evaluation of such theories via hypothesis testing is found in some sub-disciplines of computer science, including HCI, empirical software engineering, and web science. Where our department is notable, and perhaps unique, is in its *application of ergalics across computer science*, including those sub-disciplines concerned with specific software systems artifacts: databases, networks, multimedia systems, operating systems, and robotics.

All of the ergalics projects in our department (an extended list and discussion can be found at `http://www.cs.arizona.edu/projects/focal/ergalics`) share several important characteristics. They propose predictive causal models about members of an identified class of computational tool and they subject those predictive models to hypothesis testing. They thus strive to articulate fundamental properties or fundamental understanding about the behavior of those tools or of the nature of interaction with users of those tools. They all embrace empirical generalization.

The focal problem embodied by ergalics is concise:

> *What are the predictive causal models that underlie computational tools, the use of such tools, and computation itself?*

Why is this important? Succinctly: a predictive model can be tested by comparing its predictions to what we observe in experiments, lending credence to the model, thus helping us uncover causal relationships that enable *control* of the computational tool, which then enables *improvement* (in terms of performance, functionality, reliability, and other engineering goals) as well as possibly *mathematical insights* in the form of new theorems. Prediction, along with explanation, yields understanding.

This focal problem is a general form of more specific problems whose detailed answers will help advance the power, utility, and ease of use of specific computational tools and will help us understand the potential and limitations of those tools and of computation itself. The following are just some of the specific research problems being actively addressed by faculty in our department.

- How can the performance of a program executing at an arbitrary CPU frequency be accurately predicted from data collected from a single program run at a single CPU frequency?
- Does adding a physical operator for an algebraic operation always improve the effectiveness of query optimization, or is there a limit to the number of operators that can be practically accommodated?
- How could automatic video analysis of educational lectures and talks predict slide changes based on detected camera events?
- How can the impact of various components and composition of sophisticated online tutoring systems on the efficacy of learning by students of skills encompassed by stated objectives be predicted?
- How can the evolution of the topology of the Internet be predicted from low-level data, such as BGP logs, traceroute, and Internet registries?

The overarching focal problem and these specific problems that share its goal of predictive causal models cannot be answered through mathematical theorems, for we are nowhere near understanding most complex computational tools to the degree required to state and prove such theorems. These pressing problems also cannot be addressed through the building of engineering artifacts, for that activity cannot address problems about an entire *class* of computational tool. Rather, addressing this focal problem requires a new perspective, in addition to the mathematical and engineering perspectives: that of science.

Ergalics is a focal problem and an associated methodology. It is not a discipline; rather, it transcends and embraces most sub-disciplines of computer science. There are eighteen specific research projects within our department that use this methodology, studying many classes of computational tools (database ergalics, instructional ergalics, high-performance computing, multimedia ergalics, network ergalics, operating systems ergalics, and robotic ergalics), as well as more foundational efforts (algorithmic ergalics, automated ergalics, cognitive ergalics, ergalics metrology, and philosophy of ergalics). These projects involve thirteen tenure-track and research faculty and eight collaborating faculty from other departments within the University of Arizona as well as at other institutions, three industrial collaborators, one postdoc, fully two dozen current graduate students and half a dozen current undergraduate students, supported by a dozen current grants.

Understanding Human Behavior

By IAN FASEL

The research in my group spans the fields of cognitive science, artificial intelligence, machine perception, and robotics. Our goal is to learn about human cognition by developing computational systems that allow us to analyze and understand human behavior (particularly social behavior) collected from several different areas of the behavioral sciences, while in parallel developing robots that have to solve some of the same real-world problems that the human brain has to solve. To date, this has included real-time detection of objects, faces, and facial expressions, real-time classification of non-verbal speech, vocal tract gestures, and audio scenes, and active learning and perception in several other sensory domains such as haptic feedback for robots covered with flexible, touch-sensitive skin. Most of this work involves fully unsupervised learning or "weakly-supervised" learning, in which non-expert humans only provide a little, potentially ambiguous, instruction to the computer.

A current focus in the Arizona Robotics Research Group (ARRG) is "active learning" agents, i.e., agents that have to make intelligent decisions about how to move about space and orient their sensors over time in order learn about their environment as quickly as possible. Not only are our robots "born" not knowing where or what the objects in the room are, but they often don't even know the specific response characteristics of their own sensors! These all must be learned from experience. Thus for these robots, sensing is not a passive process, in which someone presents a stimuli which must then be classified, rather it is an active information gathering process that requires intelligent decision making. Key to this work is the ability to learn and reason about how moving one's own body might change one's perceptions, and how different types of objects might behave when acted upon. Visual, haptic, and other sensory information are all combined to develop much richer representations of concepts than just whether a patch of pixels looks like a face, a car, a bike, or some other object. Ultimately our aim is for agents to be able to learn and truly understand a primitive set of language-like concepts about space and physical properties that can serve as the basis for natural interaction with humans and rapid adaptation to new environments.

Another line of work involves analysis of the human articulatory apparatus during speech. Using a variety of sensors in addition to audio, including ultrasound, electroglottograpy, nasal airflow, and video, our aim is to develop a more complete understanding of the process by which humans produce sounds of all types, including language and music, and how we understand them. This work has broad potential applications, not only in phonology and linguistics, but also in speech recognition, language documentation, language learning, speech therapy, and music. This work is in collaboration with Diana Archangeli and Jeff Berry, in the UofA Department of Linguistics.

Yet another area of research in the lab is in human-computer and human-robot interaction. Several projects fall under this umbrella, including a large distributed multi-player game scenario, several different "Wizard of Oz" experiments, and even work on touch-sensitive tablets in which we recognize users' identity and specific touch gestures in real-time. Several collaborators in cognitive science, psychology, and economics are involved in this work. Despite the apparent diversity in these settings, they all share a common methodology of applying modern machine learning methods to extract deeper, hidden information about human states from temporal sequences of human biometric data. Together, our behavioral science and robotics goals may be seen as two, convergent approaches to understanding intelligence, by respectively analyzing how nature's best example of intelligent agents — namely, people — behave, and by formulating hypotheses for how to produce intelligent behavior in real-world autonomous agents. It is an exciting endeavor!

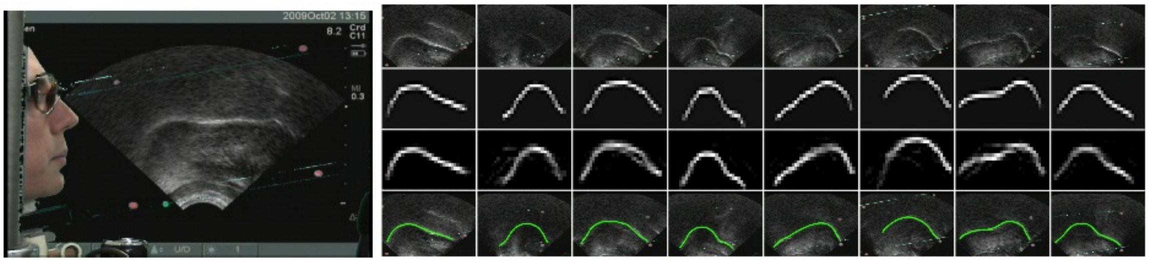

Spring and Summer 2010 Graduates

B.Sc. Computer Science

Barbara Mosier Anderson
Ethan Angus Anderson
Jose Benjie Degala Arandia
Brian Duane Bilbo
Joseph Brockert
John Caldwell Cayce
Nguyen K. Chu
Mathew Erick Cucuzza
Ian M. Demoulin
Shawn T. Fenn
Shawn Kristofor Francis
Timothy Andrew Gaddis
Erick Spock Gafni
Seth R. Gilchrist
Dustin Alexander Helak
Alex Jay Henniges
Aditya Aditama Herlambang
Wook Tae Jung
Janibek Kadraliyev
Andrew T. Kaigler
Nicolas Merrell Mardian
Daniel T. Mathis
John William McClish
Dustin R. Melich
Kenneth Steven Molnar
Amir R. Muntasser
Jonathan Scott Nation
Jude Christopher Nelson
Thuy Tuong Nguyen
Kurtis J. Norwood
Drew R. Olson
Michael Paul Ornelas
Min-sock Park
Daniel Pavel Parobek
Sean Sebastian Peterson Schnell
Scott D. Peyton
Duong Thanh Pham
Ted J. Phillips
Adam Robertson
Keven D. Schneider
Luke Monroe Shellhorn
Robert Paul Sidur
William N. Snavely
Rovert Gabriel Soimaru
Richard Thomas Solsten III
Arthur Stillwell
Joseph S. Thomas
Preston Tylor Tilus
Zhenisbek Tokabayev
Nathan Venet
Christopher Brian Waller
Kendra S. Walworth

M.Sc. Computer Science

Weitao Chen
Frederico Miguel Cirett Galan
Matthew Garst Cleveland
Ran Duan
Nithya Krishnamoorthy
Russell Lee Lewis
Jing Li
Arun Madhavan
Kyu Han Nam
Juan Antonio Raygoza Garay
Ian Ryan
Raquel Torres Peralta
Qiyam Tung
Cheng Yi

Ph.D. Computer Science

Igor Crk
Scott Morris
Joseph Schlecht
Wesley Kerr

Context-Aware Resource Management

By IGOR CRK

The demand for performance and resources that is placed on the system is largely dictated by user interactions in interactive environments. Understanding user interactions can provide valuable

information about which resources will be needed ahead of time, leading to performance optimizations like resource allocations for applications that can utilize a given resource more efficiently, and transitioning devices to a more appropriate energy performance state before the demand arrives. This work addresses the challenge of designing energy efficient systems by examining the role of user interaction in scheduling resources to adequately accommodate user demand. It is shown that system performance can be tailored to a user's pattern of interaction, resulting in applications ranging from reducing interaction time, reducing delays associated with hard disk energy management, and managing energy in wireless network cards.

Digital Trails

By SCOTT MORRIS

This dissertation presents a set of methods for collecting, improving and processing digital GPS trails, laying the ground work for the science of trails. We first present a solution to the GPS-network problem, which

This dissertation presents a set of methods for collecting, improving and processing digital GPS trails, laying the ground work for the science of trails. We first present a solution to the GPS-network problem, which determines the salient trails and structure of a trail network from a set of GPS tracklogs. A set of tracks through a GPS trail network further presents the opportunity to model and understand trail user behavior. We present the K-history model, a probabilistic method for understanding and simulating trail user decisions based on GPS data. With collections of GPS trail data we can begin to learn what trails look like in aerial images. We present a statistical learning approach for automatically extracting trail data from aerial imagery, using GPS data to train our model. These methods present further possibilities for the study of trails and trail user behavior, resulting in increased opportunity for the outdoors lover, and more informed management of our natural areas.

Learning 3-D Models of Object Structure From Images

By JOSEPH SCHLECHT

Recognizing objects in images is an effortless task for most people. Automating this task with computers, however, presents a difficult challenge attributable to large variations in object appearance, shape, and pose. The problem is further compounded by ambiguity from projecting 3D objects into a 2D image. In this thesis we present an approach to resolve these issues by modeling object structure with a collection of connected 3D geometric primitives and a separate model for the camera. From sets of images we simultaneously learn a generative, statistical model for the object representation and parameters of the imaging system. We explore our approach in the context of microscopic images of biological structure and single view images of man-made objects composed of block-like parts, such as furniture. We express detected features from both domains as statistically generated by an image likelihood conditioned on models for the object structure and imaging system. Our results demonstrate that we can infer both 3D object and camera parameters simultaneously from images, and that doing so improves understanding of structure in images.

Recognizing Behaviors and the Intentional State of the Participants

By WESLEY KERR

Psychological research has demonstrated that subjects shown animations consisting of nothing more than simple geometric shapes perceive the shapes as being alive, having goals and intentions, and even engaging in social activities such as chasing and evading one another. While the subjects could not directly perceive affective state, motor commands, or the beliefs and intentions of the actors in the animations, they still used intentional language to describe the moving shapes.

In this dissertation, we present representations and algorithms that enable an artificial agent to correctly recognize other agents' activities by observing their behavior. In addition, we demonstrate that if the artificial agent learns about the activities through participation, where it has access to its own internal affective state, motor commands, etc., it can then infer the unobservable affective state of other agents.

Departmental Awards

Graduate Student Research Award
Barry Rountree

Graduate Teaching Assistant Award
Rui Zhang

Student Outstanding Service Award
Wesley Kerr

Excellence in Undergraduate Research Award
Alex Henniges

Outstanding Senior Award (Fall)
Loren Chea

Outstanding Senior Award (Spring)
Jonathan Nation

Staff Outstanding Service Award
Tom Lowry (also CoSSAC Award)

Faculty Impact Award:
Chris Gniady

Faculty Outstanding Service Award
David Lowenthal

Academic Services

By CHERYL CRADDOCK

Fall 2010 saw Academic Services settling in following a time of transition. After many years of dedicated service to the students and faculty in Computer Science, Rebecca Mitchell, Rhonda Leiva, and Lupe Jacobo have moved on to new adventures.

Lupe is still in Gould Simpson, but has advanced to the role of special assistant to SISTA Director Paul Cohen. In January 2010, Rebecca accepted a position with the Arizona State University School of Social Work located in Tucson, Arizona. She is the Academic Advisor for the Tucson program. Rhonda has moved to Connecticut to be with her family, and has taken over the reins as an Assistant Director for Student Affairs at the Yale School of Medicine.

We're fortunate to have three experienced and energetic people stepping into the Academic Services office. Like many other units on campus, Academic Services faces the challenge of streamlining its operations while remaining flexible and effective in tending to student and academic program needs. Our new Academic Services team brings to the table wide-ranging backgrounds and talents, and aims to maintain an engaging and friendly environment that serves the students, staff and faculty of the department.

Who are we?

- In August, we welcomed Holly Brown to the position of Administrative Secretary. She worked for the Student Union as the Gallery Curator for the Student Union arts gallery prior to joining us. She hails from Pennsylvania, and came to Tucson to earn her Master of Arts in Art and Visual Culture Education. She volunteers at the Museum of Contemporary Art, and is starting a garden and welcomes advice!
- In May, Cheryl Craddock joined us in the position Program Coordinator, Senior, focusing primarily on graduate student services. Cheryl has wide-ranging experience at the UA, working and providing leadership in academic advising, program coordination and development, teaching, and research, most recently at the School of Natural Resources where she also earned a PhD in Natural Resources with an emphasis in Wildlife Conservation and Management (which helps explains the knickknacks on her windowsill).
- In August, Heather Jepsen joined the team as Program Coordinator, and is focusing mostly on the undergraduate programs. She has worked extensively at the UA, in Agriculture and Resources Economics, Agricultural Education, and as part of the Mosaic project. She has a Master's degree in Agricultural Education, and is the proud owner of two tiny dogs and one giant rabbit.

Gifts and Sponsorship

Gifts

Lockheed Martin
Cloakware, Inc.
Componentree
Microsoft
IBM

Thank you for your generosity and support this year!

Supporters

National Science Foundation
Department of Education
Defense Advanced Research Projects Agency
Air Force Office of Scientific Research
Department of Homeland Security
Raytheon/BBN Technologies
SRI International
UA Technology Research Infrastructure Fund
Computing Research Association

Industrial Partners

Lockheed Martin
Google
Rincon Research Corporation
Jigsaw Technologies