


WE ARE COMPUTER SCIENCE & ENGINEERING

UNIVERSITY OF CALIFORNIA SAN DIEGO



UC San Diego

JACOBS SCHOOL OF ENGINEERING
Computer Science and Engineering

9500 GILMAN DRIVE, DEPT 0404
LA JOLLA, CA 92093-0404

CSE.UCSD.EDU

1 Letter From the Chair

2 CSE By the Numbers

4 Awards & Honors

Research

- 8 Robotics
- 9 Embedded Systems
- 10 Artificial Intelligence
- 14 Databases
- 15 Theory
- 16 Systems & Networking
- 18 Security & Cryptography
- 20 Human-Computer Interaction
- 22 Computational Biology
- 24 Graphics & Vision
- 26 Programming Systems
- 28 Architecture
- 29 VLSI/CAD

30 Alumni

32 Undergraduate Education

34 Graduate Education

36 Faculty

Letter from the Chair

▶ **IT IS AN AMAZING TIME** to be a computer scientist, as the discipline has moved from the fledgling, upstart specialty it was several decades ago to one of the primary fields in engineering and academia as a whole. We see this not only in the numbers that are flocking to the discipline, but also in the transformative impact it is having on the world today. Computer science has changed the way we learn, the way we work, the way we play, the way we travel, the way we communicate – in short, the way we live.

The pages of this brochure are filled with stories of all these exciting transformations. But these stories show something else as well: there is no better place to share in this excitement than the Department of Computer Science and Engineering at UC San Diego.

We are one of the top academic departments in the world, with strong research groups spanning all the major areas of computer science, including computer architecture, bioinformatics, databases, embedded systems, graphics and vision, machine learning, natural language processing, networking, programming languages, robotics, security and cryptography, software engineering, and theoretical computer science. Our faculty are internationally recognized experts and include many IEEE and ACM Fellows. They are also committed educators who have won a slew of campus-wide and national awards for teaching.

As the field has grown, so has our department along with it. We have added faculty at an aggressive rate, with 21 new hires in the last 4 years, and welcomed our largest incoming class of PhD students. We have completed a major renovation and addition to our building, and many of our faculty helped to launch the new Halicioğlu Data Science Institute at UCSD. With our powerhouse reputation for research, and our proximity to some of the best coastline of Southern California, we continually attract many of the best researchers in the world – faculty and students alike.

Please come visit the Department of Computer Science and Engineering at UCSD and see all the excitement for yourself. While you're here, you'll also discover why San Diego is one of the most desirable places to live in the world. But in the meantime, please take a look at our brochure, and let it introduce you to the breadth of our research, the excellence of our teaching, and the impact of our students and faculty on the world around them.



CSE CHAIR DEAN TULLSEN
University of California San Diego



We have added faculty at an aggressive rate, with 21 new hires in the last 4 years, and welcomed our largest incoming class of PhD students."



COMPUTER SCIENCE & ENGINEERING

CSE

UNIVERSITY OF CALIFORNIA SAN DIEGO

CSE By the Numbers

250

PHD

520

MASTERS

2,019

UNDERGRADUATES

9,957

ALUMNI

45

STAFF

71

FACULTY

AWARDS & HONORS

All researchers pursue new understandings, insights or capabilities in the hope that their work will change the field and, ultimately, the world. But it is only with the passage of time that the field can fairly assess the impact of the work. With this retrospective view, some ideas and accomplishments are so clearly influential that the community stands up and recognizes their merit publicly.

Time-tested research

► Many leading computer science conferences have instituted an annual Test of Time award celebrating individual papers that have made a clear and unique contribution over the preceding decade. Between 2017 and 2019, research published from CSE has received this honor seven times: **George Porter** won the **USENIX NSDI Test of Time Award** in 2017 for “*X-Trace: A Pervasive Network Tracing Framework*,” and **Geoffrey Voelker** won it in 2019 for his work on “*Sora: High Performance Software Radio Using General Purpose Multicore Processors*.” **Voelker** and **Stefan Savage** won the **2017 USENIX Security Test of Time Award** for their 2000 paper “*Inferring Internet Denial-of-Service Activity*.” During the same period, **Alin Deutsch**’s paper “*The Chase Revisited*” was recognized with the **2018 PODS Test of Time Award**, and he and **Victor Vianu** also were honored with the **2019 ICDT Test of Time Award** for their paper “*Automatic Verification of Data-centric Business Processes*.” Finally, faculty alumni **Amin Vahdat** and **Hovav Shacham** (and their students) received **Test of Time awards** from ACM SIGCOMM and ACM CCS, respectively, for their work published while at UC San Diego 10 years ago.

Leaders in their field

► Each sub-community inside computer science identifies and recognizes its leaders and leaders-to-be via area-specific awards. Among these, in 2018, the IEEE Technical Committee on Computer Architecture honored **Hadi Esmaeilzadeh** with the **Young Computer Architect Award** for his work on novel computer architectures designed to efficiently support machine learning applications. The same year, **Ranjit Jhala** received the **ACM SIGPLAN Robin Milner Young Researcher Award** for his “foundational and impactful contributions to programming languages,” including his work on BLAST and refinement types. Finally, rewarding a long career of research leadership, the International Society for Computational Biology (ISCB) named **Pavel Pevzner** as the recipient of its **2017 Senior Scientist Award**, recognizing researchers “who have made major contributions to the field of computational biology.”

ACM & IEEE Fellows

► The Association for Computing Machinery (ACM) is the world’s largest educational and scientific computing society, and each year it recognizes a few dozen members as Fellows to honor their “achievements in computer science and information technology.” Fewer than one percent of ACM’s 100,000 computing professionals are ever named Fellows. Over 20 faculty in CSE have received this honor, and joining the list in 2017-2019 are **Rajesh Gupta** (for “contributions in design of embedded systems and hardware-software co-design”), **Ravi Ramamoorthi** (for “contributions to computer graphics rendering and physics-based computer vision”), **Alexander Vardy** (for “contributions to the theory and practice of error-correcting codes and their study in complexity theory”), **Geoffrey Voelker** (for “contributions to empirical measure-

ment and analysis in systems, networking and security”) and **Alex Snoeren** (for “innovative approaches to measuring, managing and detecting network traffic”). As well, **Tajana Šimunić Rosing** and **Ravi Ramamoorthi**, were also recently elected Fellows of the Institute of Electrical and Electronics Engineers (IEEE), the world’s largest technical professional organization. Rosing was honored for her “contributions to power and reliability management of systems-on-chip,” and Ramamoorthi for his “contributions to foundations of computer graphics and computer vision.” They join the 12 other CSE faculty who have been named IEEE fellows for their body of work.

Recognizing a body of work

► Finally, two of our faculty were recently recognized for the lasting contributions of their work. **Andrew Kahng** was named the winner of the **2019 Ho-Am Prize for Engineering** (the most prestigious engineering award given in Korea). The Ho-Am Foundation recognized him for introducing the concept of design for manufacturability (DFM) and for design technologies that “have become the standard ... around the world.” In the same year, **Rajesh Gupta** was named the recipient of the **IEEE Computer Society’s McDowell Award** for his “seminal contributions in the design and implementation of microelectronic systems-on-chip and cyberphysical systems.” The **McDowell Award** is presented for broad contributions to the field of computing, and previous recipients include **Lawrence Roberts** (for packet switching), **Tim Berners-Lee** and **Marc Andreessen** (for the World Wide Web and Web browser) and UCSD Chancellor **Pradeep K. Khosla** (for reconfigurable real-time systems).

MACARTHUR FELLOW



THE LAST FIVE YEARS have been good to CSE Professor Stefan Savage. In 2013, he was awarded the ACM SIGOPS Mark Weiser Award for creativity in systems research; in 2015, he received the ACM Prize in Computing for work that “taught us to view attacks and attackers as elements of an integrated technological, societal and economic system” and in the summer of 2017, he and fellow CSE Professor Geoffrey Voelker received the USENIX Security Test of Time Award for their research on denial of service attacks.

Then, in the fall of 2017, Savage started receiving calls on a daily basis, sometimes several times a day, from a phone number with a Chicago area code. The caller didn’t leave a voicemail and, having been targeted by scammers before, Savage never answers a number he doesn’t recognize. Finally, Savage looked up the phone number: the call came from the MacArthur Foundation’s headquarters. A few calls later, after proving his identity, he found out he was a 2017 MacArthur Fellow. Better known as the MacArthur “genius” award, the prestigious, no-strings attached, five-year fellowship awards a total of \$625,000 to each recipient. Savage was recognized for “identifying and addressing the technological, economic and social vulnerabilities underlying Internet security challenges and cybercrime.”

Indeed, Savage has taken an expansive view of what it means to study cybersecurity. He and his colleagues bring together computer science and the social sciences, taking into account not just technology, but also economics, policy and regulations. “Fighting cyber threats requires more than just understanding technologies and the risks they’re associated with; it requires understanding human nature,” Savage says. “At its heart, cybersecurity is a human issue. It’s about conflict, and computers are merely the medium where this conflict takes place.”

Savage’s research has also been broad in scope. He, with students and collaborators, has remotely taken over cars to reveal security vulnerabilities, tracked criminal financial transactions responsible for funding online scams and botnets around the world, found ways to pinpoint the source of attacks that cripple the

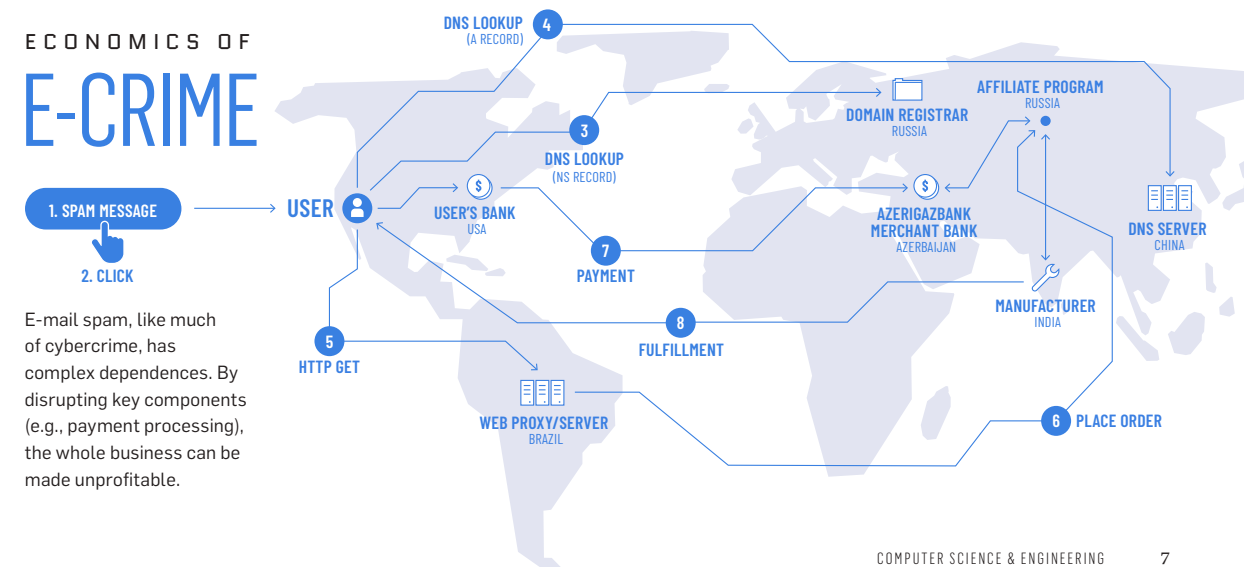



UCSD alumni **Ian Foster** and **Karl Koscher** debugging a car hacking demo

Internet and to identify previously unknown data breaches. As one of his award nominators wrote, “Our inability to select a single ‘greatest hit’ does not make The Beatles a lesser band; rather, we recognize that any one of their better songs would have been sufficient to catapult a lesser band into the Rock & Roll Hall of Fame. Ditto for Savage.”

Finally, and perhaps the most unusual, Savage’s work has routinely led to more than just research papers, but also impact in the real world – changes in how automobiles are now designed and built, and in how the financial system is being used against cybercriminals.

Today, Savage, his colleagues and their students are working hard to discover their next big research hits in new domains ranging from aviation to health care, usability and security metrics.





Artificial Intelligence

► **MOST PARENTS CAN LOOK AT A CHILD'S FACE** and know when they are in pain. One might expect doctors, whose job is to treat human suffering, to be even more adept at gauging pain from facial expressions. In fact, the opposite is true: study after study has shown that physicians are worse at this than the general population.

CSE Professor Laurel Riek and her students have now taken this finding one step further. They have shown that clinicians also perform worse than laypersons at gauging pain from the expressions of humanoid robots and virtual avatars. Their findings have implications for how clinicians should be trained to better recognize and treat pain.

One effective way to train clinicians is to use patient simulators. Unfortunately, existing patient simulators do not make expressive faces, and commercial expressive robots (such as Hanson Robotics' platforms, costing \$50K-\$100K) are either too expensive or impractical to be used in clinical settings.

Riek is hoping to fill this gap. Her team is building a low-cost robotic patient simulator (RPS) capable of facial expressions, with a target price tag of \$10K. In addition, they are designing virtual avatars to train clinicians in low-resource environments, such as developing countries. This work is supported by Riek's NSF CAREER Award.

The first step of this research was to develop computer vision algorithms that analyzed the facial expressions of people in pain. The algorithms tracked the movement of more than 60 landmarks on the human face, modeling them as vertices of a dense triangular mesh. The results were then applied to synthesize expressions on an expressive RPS and a digital avatar.

"From the models we've built," Riek says, "we're able to synthesize conditions such as pain, stroke, and facial paralysis."

In one study, the researchers showed videos of the virtual avatar and robot to 102 volunteers. The volunteers — half clinicians, half laypersons — had to decide which emotion they were seeing in the videos. Overall, laypersons correctly identified expressions 67 percent of the time versus 51.6 percent for clinicians. Both groups of volunteers did better with the digital avatar than with the robot. The paper was published at the ACM/IEEE International Conference on Human-Robot Interaction.

In a follow-up study, Riek's team modeled facial paralysis in Bell's palsy, a condition that is sometimes misdiagnosed by clinicians. In experiments on a virtual patient simulator, their synthesized expressions were shown to be realistic and comparable to those of real patients. This paper was published at the IEEE Conference on Automatic Face and Gesture Recognition.

Riek hopes that repeated training on and exposure to an expressive RPS will help clinicians improve their ability to recognize and rate pain in human patients, as well as to better treat patients with facial paralysis. They will incorporate these findings in the next stages of their work. Riek also envisions an RPS that could be outfitted with diverse faces that convey different ages, genders, and ethnicities "There are extreme health disparities in this country, with people at social, economic, and environmental disadvantages receiving poorer care," she says. "Our goal is to improve treatment for everyone."



Laurel Riek is a roboticist with interests in human-robot teaming, computer vision and healthcare engineering.


HOW IT WORKS:



1. Source Video
Camera captures the operator's face in real time, or the operator selects a pre-recorded video.



2. Facial Features
Frame by frame, a face tracker extracts facial feature points of the operator's face.



3. Feature Masking
Operator selects a pre-built mask of a patient pathology (e.g., Bell's palsy, stroke) based on models built by Riek's team. The system then applies the scaling parameters of the mask to the extracted facial feature points.



4. Synthesis on Patient Simulator
Then the masked features are synthesized on either a virtual or robotic patient simulator in real time.

Memory-centric Computing

► **MODERN COMPUTERS ARE PROCESSOR-CENTRIC** — data is stored in external memories and transferred to individual computing cores for processing, and then the results are transferred back. This arrangement has served for decades, but it imposes implicit performance and power bottlenecks on emerging data-intensive applications.



If things work out, we could support real-time machine learning on small devices.... Instead of sending data to the cloud, as we do today, we could do all the training or processing in a sensor, a watch or cell phone."

Instead, one alternative is to combine memory and processor into a single unified device. **Professor Tajana Šimunić Rosing** and her students are working with the DARPA/SRC-funded Center for Research on Intelligent Storage and Processing-in-Memory (CRISP) to make this happen. "We're designing new systems that place processing inside nonvolatile memories. The computation happens right where the data is stored," says Rosing. "This has a huge impact on applications with large datasets."

Her team is not only focused on improving the flow of data, they have their sights set on vastly improving computational efficiency as well. Instead of using processors that perform arithmetic on numbers, Rosing is advancing a brain-inspired architecture that processes data in hypervectors, large chunks of data representing multidimensional features. This pattern-matching orientation for computing naturally maps onto large classes of machine learning problems. By designing smart nonvolatile memories capable of processing hypervectors natively, early results suggest this approach can improve computing efficiency by multiple orders of magnitude.

"If things work out, we could support real-time machine learning on small devices," says Rosing. "Instead of sending data to the cloud, as we do today, we could do all the training or processing in a sensor, a watch or cell phone. What today requires hundreds of thousands of remote servers could all happen in your hand."




Tajana Šimunić Rosing works at the intersection of software and hardware design. Her current research interests are in the area of cyber-physical systems, emerging computing and system-level energy efficiency.

FRUIT FLY SEARCH ENGINES

Fruit flies rely on a collection of neurons to distinguish new odors from those that they have already encountered. In computer science, this important task is called novelty detection, and it is used to identify duplicates or anomalies in large data sets. It turns out that computers and fruit flies identify novel information in similar ways.

This surprising connection is the latest result from an interdisciplinary collaboration between researchers in CSE and the Salk Institute's Integrative Biology Laboratory. Their findings were reported in a recent issue of the *Proceedings of the National Academy of Sciences (PNAS)*. CSE co-authors include **Professor Sanjoy Dasgupta** and graduate student **Timothy Sheehan**, who are studying the problem of novelty detection from the perspective of machine learning.

When a search engine such as Google crawls the Web, it needs to know whether a website it comes across has previously been indexed; otherwise it may waste time indexing the same site again. The problem is that there are trillions of websites on the Web, and it is prohibitively expensive to store all of



Sanjoy Dasgupta
develops algorithms
for unsupervised and
interactive learning.

them in memory. In the 1970s, long before the Web was invented, Howard Bloom at MIT devised a filter to solve this problem: for each item in a large database, the filter stores a compact “fingerprint” that requires only a few bits of memory. By checking whether the same fingerprint appears twice in the database, a system can quickly determine whether an item is a duplicate or something new.

It seems that fruit flies do something similar. Fruit flies are well-known to change their behavior in response to novel odors. A region in the fly brain, called the mushroom body, contains a collection of neurons that processes olfactory information. When a new odor is experienced, these neurons broadcast a “novelty alert” signal so the fly knows this odor is new and worth investigating. This novelty signal is generated using an odor fingerprint analogous to the data fingerprint in the Bloom filter.

Along with lead author Saket Navlakha of the Salk Institute, Dasgupta analyzed the fly’s novelty neural circuit and found that the fruit fly introduced a couple of twists on the traditional Bloom filter. The first is to test for approximate rather than exact matches so that the fly can detect odors that are similar, but not strictly identical, to ones it has previously encountered. The second is to encode temporal information so that an odor is more likely to be registered as novel if it has not been smelled in a long time.

Based on the fruit fly’s variant of the Bloom filter, the team created a new algorithmic framework to predict the novelty responses of the fly’s mushroom body neurons. Testing this same framework on several data sets in machine learning, Dasgupta and his colleagues also found that it improved on other known methods for novelty detection.

“A fundamental challenge in machine learning is finding representations of data that are suitable for a wide range of tasks,” says Dasgupta. “The fly’s olfactory system shows us one simple and ingenious way of doing this.”



FRUIT FLY FACTS

01_ Nearest neighbor search can be an effective strategy for pattern recognition; the main question is how to do it efficiently.

02_ In general, it’s much easier to search through sparse representations of data. One template for sparseness is suggested by the olfactory system of the fruit fly.

03_ The fly’s sensory inputs are randomly projected into a large number of dimensions; then only a few top values are retained to create a sparse representation.

04_ The resulting representation can be hashed for nearest neighbor search. The overall approach, combining sparse random projections and hashing, is being studied for practical applications.

Protecting Data Privacy

► **CSE PROFESSOR KAMALIKA CHAUDHURI** has a big question she's trying to answer: "How do you build predictors out of sensitive data that preserve people's privacy?"

As more personal data is collected and analyzed to solve problems, it becomes increasingly important – and difficult – to keep that data secure. Advances in computing technology are delivering faster processors and more efficient algorithms. But as more information is brought into the system, privacy only gets harder.

"I can get genomes from 100 people, then build a classifier to predict whether a person is prone to certain diseases," says Chaudhuri. "But if my classifier overfits to the genomes in the training data, someone may be able to deduce who was in the training dataset to begin with."

Specificity is important to understand the data, but it can also destroy privacy. When presenting on her research, Chaudhuri is often asked why anonymity is not enough. She uses herself as an example.



Kamalika Chaudhuri researches the foundations of trustworthy machine learning.

"Let's say UC San Diego decides to release a salary table, including information about department, gender and ethnicity," says Chaudhuri. "Department is CSE, gender is female, ethnicity is South Asian. The moment you see these fields, you know it's me because I'm the only one who fits this description."

This becomes particularly tricky when only a few people are involved. Larger datasets are more private, but they are also more expensive to produce.

"We're trying to understand how to design more complex systems that will preserve privacy," says Chaudhuri. "Usually the way this is done is by adding a bit of statistical noise to hide individual effects."

But this leads to new questions: What kinds of noise, how much noise, and how do you add it to the dataset?

These questions have led to a wave of pioneering research. Chaudhuri and colleagues have developed privacy-preserving frameworks for some of the most popular methodologies in machine learning, including large-margin classification, logistic regression, and Bayesian data analysis. They have also explored privacy issues for data that are correlated across time or via connections in a social network.

Networked data arises often in patient care. To better address privacy issues in this space, Chaudhuri has teamed up with researchers in the School of Medicine at UCSD.

There are other areas where private information can be gleaned from complex data. Location data, for example, can be used to determine where a person is going. Chaudhuri feels these issues often fly under the radar, making little or no impact on public perceptions. She wants to change that.

"It's important to have public awareness," she says, "so that when policies are on the table, the public can make an informed decision."



Neural Nets on the Fashion Runway



Julian McAuley uses large volumes of online activity data to develop predictive models of human behavior.

► **MOST OF US WILL NEVER HAVE OUR OWN CELEBRITY FASHION LINES**, but advances in AI and machine learning may soon deliver the same personalized experience. Working together, computer scientists from UCSD and Adobe Research have shown how neural networks can learn to create custom apparel designs based on a buyer's purchase history. The research is part of a long line of work on next-generation recommender systems, led by **Professor Julian McAuley** in the department's machine learning group.

"What's new here is that our models can be used generatively," says computer science **Ph.D. student Wang-Cheng Kang**, who is first author on the paper presenting these results at the International Conference on Data Mining. "In other words, given a user and a product category, we can generate new clothing items that are most consistent with an individual user's personal taste."

Initially, the researchers focused on using visual information to improve recommendations from an existing collection of products. To do this, they trained a recommender system based on purchase histories jointly with a convolutional neural net that analyzed raw images. In this way the neural net was able to learn "fashion-aware" image representations directly from pixel inputs.

From there, the team used a generative adversarial network (GAN) to learn the distribution of fashion images that maximize an individual user's preferences. The resulting system can not only suggest items to buy from an existing collection of products; it can also be used to modify existing items or to generate novel fashion items with specially tailored designs. The ultimate goal is to help retailers and apparel makers sell bespoke clothing to a worldwide customer base.

"Building effective recommender systems for domains like fashion is challenging due to the high level of subjectivity and the semantic complexity of the features involved," the researchers say. "But this represents a first step toward building systems that go beyond recommending existing items from a product corpus, to suggesting styles and helping to design new products."

The use of AI in the fashion industry is still in its infancy. The quality of the algorithmically designed new clothes remains rudimentary at best. But two of the world's largest online retailers – Amazon and China's Alibaba – are already working with AI tools such as GANs. While fashion can be notoriously unpredictable, one thing seems certain: the role of AI in this industry is set to grow.

Actual

Algorithmic Alterations



Clothing designs can be algorithmically altered to produce designs that better match user preferences. Above, the same catalogue item (left) is modified to produce similar synthetic items (right) that more closely match the preferences of two different users.

Integrating Data Management and Machine Learning

► **DATA SCIENCE IS BLENDING STATISTICAL AND MACHINE LEARNING (ML)** with systems to analyze large and complex data sets. But, according to **CSE Professor Arun Kumar**, the gaps between ML and data systems-oriented thinking are driving new computational challenges and opportunities.

Kumar's research bridges both of these worlds.

"Surveys of data scientists show that a vast majority spend almost 90 percent of their time on painstakingly preparing the data and configuring ML to work — not on ML algorithmics," notes Kumar. This highlights the disconnect between ML and data processing needs.

By abstracting the steps in the ML-based analytics process, Kumar leverages fundamental computational and statistical properties to build better ML systems, making it faster and easier to analyze data.

"A major example is ML on multi-table datasets," says Kumar. "For instance, a recommender system, such as Netflix, has separate tables for users, movies, and ratings. To predict future ratings, data scientists combine all these tables into one massive table using a database operation called joins."

Unfortunately, joins introduce redundancy in the data and ML computations, wasting storage, memory, and runtime.



Surveys of data scientists show that a vast majority spend almost 90 percent of their time on painstakingly preparing the data and configuring ML to work — not on ML algorithmics."

Arun Kumar works at the intersection of data management, machine learning and systems. His current focus is on improving the efficiency, scalability and usability of machine learning-based analytics over large and diverse forms of data.



"If a movie has a million ratings, its record is repeated a million times after the join," he explains. "Data access controls within organizations can also make it hard to even procure all tables for joins." Overall, joins before machine learning waste system resources and reduce productivity.

In response, Kumar introduced a technique called "avoiding joins safely," which has a surprising premise: joins are often not needed for accurate predictions. His work revealed, for the first time, the deep connection between logical database properties and the bias-variance tradeoff in ML.

"One can often ignore the movie table entirely and use the movie identifiers in the ratings table, called the foreign key, as a representative," says Kumar. "Our analysis characterized exactly how and when this change affects accuracy."

Kumar's lab distilled this idea into easy-to-understand heuristic decision rules so that data scientists can know up front when avoiding a table might significantly affect accuracy. Crucially these rules depend only on the database metadata, such as the numbers of records in the tables. This can save data scientists the hassle of procuring and joining extra tables.

Kumar and his collaborators showed their decision rules are robust and apply to many types of classifiers, including linear models, decision trees, and neural networks.

"They are also usable in any data system setting — in Python or R, and in scalable tools such as Spark, TensorFlow and others," adds Kumar. "This technique has already been adopted or explored for use by many companies, including Facebook, Google, LogicBlox, and MakeMyTrip. It signals the growing importance of integrating ML-oriented and data systems-oriented thinking."

Which image is a better search result for sky?
The answers to such comparison queries can be exploited by active learning algorithms.



Streamlining AI with Active Learning

► **ARTIFICIAL INTELLIGENCE HAS THE POTENTIAL** to address a wide range of computational problems. However, one of the bottlenecks is giving learning algorithms enough information to produce accurate models. To overcome this, some computer scientists are embracing a strategy called active learning.

"Here's the challenge," says **CSE Professor Daniel Kane**. "Suppose you want to train a computer to classify and sort images. To do this, you need a data set of labeled images, meaning that for each image, you also have a label indicating what's in it. What's more, for most applications, these ground truth labels must be collected from human annotators."

Active learning seeks to streamline this process, which can be slow and cumbersome, especially for large datasets. It gives learning algorithms the ability to query human annotators, which can potentially reduce the number of labeled examples that are needed.

"Maybe you have this large dataset, but having some examples labeled will be much more informative than others," says Kane. "An algorithm can select certain examples to be annotated, and by doing that, it can reduce the effort required of human annotators."



Daniel Kane works across multiple areas of mathematics and theoretical computer science. His work in computational statistics produced the first computationally efficient algorithms for overcoming adversarial noise in high dimensional data.


Active learning algorithms query human annotators for information about currently unlabeled examples. The simplest queries ask outright for the label — for example, whether or not an image contains a portion of the sky. Kane and colleagues analyzed the power of more subtle queries in which the annotator is asked to compare two examples. In these so-called comparison queries, the annotator might instead be shown two images and asked which one is a better search result for the word "sky."

"Using these comparison queries, you can get more efficient algorithms with fewer labels," says Kane. "The idea is that we'll have less dependence on human annotators."

In fact, Kane and colleagues showed that for the learning of half-spaces — perhaps, the simplest model of binary classification — comparison queries can offer exponential improvements over label queries. This approach has also yielded efficient algorithms for other types of problems, such as determining whether a list of numbers contains three that sum to zero.

The next challenge is to bring this active learning approach to real-world problems. "We have this working in very clean theoretical settings," says Kane, "but there's a little bit of work to relax our assumptions."

Right now further research is needed to bridge the gap between theory and practice. For Kane, an especially active theoretician in CSE, the work is never done.



Switching on the Lights

► **AT THE CORE OF TODAY'S INTERNET SERVICES** are data centers made up of tens of thousands of high-performance servers. These servers must process huge amounts of data in concert, and thus the scalability of these modern data centers is heavily dependent on the underlying communications network. Indeed, modern data center network design has been transformed over the last two decades, from traditional “scale-up” designs limited by the speed of a few core switches to modern “scale-out” networks using parallel links and complex layered topologies to distribute traffic. However, as our insatiable demand for data processing grows, even these designs are becoming a scaling bottleneck. “Modern servers each have 100Gbps links; each switch will be expected to handle 64 such ports, and there may be five or six layers of switches to get between any two servers,” says **CSE Professor George Porter**. “At that speed and scale, the heat and energy required to repeatedly identify and switch each individual packet can be enormous.”



For the dense network fabrics needed in datacenters, all-optical switch designs are faster, cheaper and use dramatically less power.”

Porter has been working on this problem with a team of faculty and students across the CSE and ECE departments — a collective expertise that spans low-level device design to high-level scheduling algorithms. Their solution has been to radically rethink data center network architecture, replacing today's complex multi-stage topologies of packet switches with designs that rely on a reconfigurable optical core. The first such switch, Helios, used electro-mechanical optical switches to dynamically set up circuits between pairs of switch ports. Such physical light paths let data flow at almost arbitrary speed, but reconfiguration was slow, and the switch could only support highly predictable traffic patterns. Mordia, its successor, used 2D MEMS mirrors to shrink reconfiguration time down to 12 microseconds and supported a far wider range of traffic demands. A third switch, REACToR, combined a high-speed optical core for bulk traffic with a slower packet switch to handle unpredicted dynamic traffic.

However, the most radical of Porter's designs is RotorNet, a purely optical switch architecture that sidesteps the need to adaptively reconfigure in response to traffic. Instead, RotorNet switches are constantly reconfiguring — light is bounced off a rapidly spinning rotor whose diffraction gratings implicitly create new switch configurations. This simple design requires no central control, no adaptivity and can efficiently scale to thousands of ports. “The abstraction of packets is key to the Internet, but there are many ways to support that abstraction,” says Porter. “For the dense network fabrics needed in datacenters, all-optical switch designs are faster, cheaper and use dramatically less power.”

Supercomputing for a Second

► **THE EVOLUTION OF COMPUTING** has been driven by changes in technology (what is possible) and economics (what is cost-effective). Early computing services were expensive and centralized, with each user billed for the resources they consumed. Today, large-scale cloud infrastructures have revitalized this centralized computing model and now allow businesses to quickly and easily scale their computing needs on demand.

However, existing cloud service interfaces are implicitly structured around the needs of businesses and not consumers. “End users care about latency, not throughput,” says **CSE Professor George Porter**. “I don't usually want to process 1,000 videos; I want to process my video 1,000 times faster.”

That goal drives a new approach Porter calls “burst parallel computing,” in which user applications are able to efficiently harness hundreds or thousands of servers for short periods of time — even fractions of a second — to dramatically improve the performance of individual user applications. Making this approach economically feasible are new cloud function execution services, such as Amazon's Lambda, Google's Function and Microsoft's Azure Functions. Originally designed to support Web microservices, these services provide a mechanism for cheaply and efficiently dispatching many small pieces of code to be executed on cloud servers. However, with appropriate system support to manage scheduling and data transfer, Porter has shown that this same mechanism can be leveraged for individual user applications.


Working together with **CSE Professor Geoffrey Voelker** and their students, Porter demonstrated for the first time this approach in a video processing framework called “Sprocket.” Sprocket is designed to offload compute-intensive tasks, such as video compression or object recognition, decompose them into thousands of distinct subtasks and dispatch these tasks (and associated data)

George Porter's interests span distributed systems, networking and data-intensive computing. His research is focused on improving the ways that networks support data-intensive operations and enabling developers to solve important problems at low cost and with low energy requirements.



I don't usually want to process 1,000 videos; I want to process my video 1,000 times faster.”

to the cloud for execution. Thus, while two hours of video might take hours to compress on an individual user's machine, Sprocket can complete the same task in a few tens of seconds at a cost under a dollar. Generalizing this vision is an ongoing research focus for the team. “Being able to use a thousand servers for a couple seconds is not something most developers have thought about,” says Voelker. To minimize this burden, they are exploring the most effective ways to structure the burst parallel processing pipeline and what kinds of interfaces and system services can best encapsulate the associated complexity.



Validating Cryptography

► **MODERN CRYPTOGRAPHIC ALGORITHMS** can offer strong security guarantees that resist attacks from even the most determined adversaries. However, the source of this strength — formal proofs of security — depend on a range of explicit and implicit assumptions. Thus, a cryptographic algorithm may be strong in its pure Platonic form, but its real-world implementation may be brittle and weak. Sadly, this brittleness does not manifest in any obvious way, and cryptographic implementations with faulty assumptions still appear to operate correctly.

“Cryptographic assumptions represent a tough engineering problem,” says **CSE Professor Nadia Heninger**. “They can be obvious to the algorithm designer, yet may seem irrelevant to the implementor.” How to identify this disconnect is a key part of Heninger’s research. “My work stands between theory and practice,” she says. “We know from experience that it is not enough to prove an algorithm secure in isolation, but we also can’t test the security of real-world crypto implementations without understanding their theoretical basis.”

Heninger’s signature approach to cryptographic testing combines mathematical cryptanalysis techniques with Internet-wide measurement and hundreds of thousands of hours of computation. In this manner she has been able to identify a wide range of subtle vulnerabilities in the cryptography we all use. For example, RSA is the world’s most widely used public-key cryptosystem, but its correctness depends on key pairs being generated randomly and independently. By crawling the Internet and then analyzing over 11 million public keys for common factors, Heninger and her collaborators showed that tens of thousands of Internet hosts used faulty random number generators that silently compromised their security.

This failure is sadly not unique, and Heninger has documented how a range of cryptographic deficiencies — broken randomness, imperfect




My work stands between theory and practice. We know from experience that it is not enough to prove an algorithm secure in isolation, but we also can’t test the security of real-world crypto implementations without understanding their theoretical basis.”

Nadia Heninger’s research focuses on applied cryptography and security, particularly cryptanalysis of public-key cryptography in practice.



parameter validation and biased key generation — can be used to break the cryptography in a range of standards, products and services. Indeed, in one well-known example, she and her colleagues showed how subtle covert changes made to the Juniper’s ScreenOS random number generator (likely by a state actor) allowed its encrypted virtual private network traffic to be decoded on demand.

“We see these kinds of subtle vulnerabilities again and again,” says Heninger. “Today, it is easy to ship a standards-compliant implementation of a popular cryptographic protocol that is also completely insecure. Ultimately,” she says, “we need a development and validation methodology that can identify these kinds of issues before we ship products.”



Scanning for Skimmers

► **PAYMENT CARDS HAVE LONG BEEN THE DOMINANT PAYMENT** mechanism for retail commerce in the United States. Credit, debit and store-specific cards are used hundreds of millions of times a day to transfer over a hundred trillion dollars each year. Unsurprisingly, this same ubiquity has attracted criminal interests, and miscreants employ a variety of means to steal card data and then make purchases against their owners’ accounts. Among the most insidious of these mechanisms is the “card skimmer,” a device that modifies an existing retail card reader to steal a copy of the card data (and PIN) as each customer swipes their card through it.

This approach is particularly prevalent in gas station pumps, which both handle large numbers of customers and are poorly monitored. Indeed, the newest gas station card skimmers are installed entirely inside the pump itself, and thus are visually undetectable to customers and store personnel alike. By staying undetected, these skimmers may continue to steal card data for weeks or months.



Aaron Schulman studies how to integrate unconventional embedded systems into computer networks, physical infrastructure and mobile devices. He is particularly interested in problems of wireless communication.

However, a team led by **CSE Professor Aaron Schulman** and CSE alumnus **Kirill Levchenko**, have recently developed a new technique for quickly detecting these covert devices based on the wireless signals they emit. Using information obtained via a partnership with the U.S. Secret Service, Schulman’s team determined that Bluetooth was commonly being used to allow criminals to retrieve the collected card data from each pump. They hypothesized that the details of how the wireless protocol was employed might produce a unique signature distinguishing such skimmers from other Bluetooth devices. Working with partners, they canvassed over 1,100 gas stations in six states and refined their technique into a smart phone app called Bluetana that can detect a likely skimmer in just a few seconds. Working with law enforcement partners, (who were able to inspect suspicious pumps and identify any skimmers present), they were able to validate Bluetana’s alerts.

In one year of operation, Bluetana has led to the discovery of 75 Bluetooth-based skimmers across six U.S. states, all of which were recovered by state and federal gas station inspectors. “We were surprised that there were so many skimmers out there that had not been discovered by other detection methods, such as manual inspections,” said Schulman. “We found two skimmers that had evaded detection for six months, and even found well-camouflaged skimmers that had been missed by manual inspection.” Bluetana is now in use by multiple enforcement agencies in several states and is estimated to have stopped millions of dollars in theft.



← Weibel and his students testing a virtual reality prototype for expert surgeons to instruct remote novices in the field.

Healthcare Meets Human-Computer Interaction



Nadir Weibel works at the intersection of ubiquitous computing, human-computer interaction, large-scale design, and healthcare.

► **HOW CAN WE USE COMPUTERS** to provide better medical care? Some ways are obvious — for example, electronic medical records and computed tomography (CT) scans. But others require more creative thinking.

That creativity is on full display in the Human-Centered and Ubiquitous Computing Lab at UC San Diego. The lab sits at the confluence of computing, healthcare, and design, and its members are developing cutting-edge technologies for new forms of patient care.

“Think about sensors, cameras, and wearable devices,” says **Nadir Weibel**, a professor in CSE and head of the lab. “How can we use these technologies in healthcare?”



We project the ultrasound image onto the patient with augmented reality using HoloLens goggles. It's like giving physicians X-ray vision."

AUGMENTED REALITY: Weibel is collaborating with physicians at UCSD, Kaiser and the U.S. Navy to solve a wide range of problems. New technologies are developed in the clinic and/or prototyped in the Simulation Training Center.

One recent project addressed a peculiar challenge of ultrasound-guided procedures. These procedures are common in medicine, but often the ultrasound screen is across the room from the patient, forcing clinicians to look away from their patient.

Weibel and his students developed a solution based on augmented reality. “My approach to health and technology is to give superpowers to people,” he says. “We project the ultrasound image onto the patient with augmented reality using HoloLens goggles. It’s like giving physicians X-ray vision.”

Another project uses augmented reality to help novice medics perform complex surgeries. This technology is being co-developed with the U.S. Navy, which can’t always put a trained surgeon in the field. By wearing virtual reality goggles, a medic in the field is able to receive assistance from a surgeon thousands of miles away. The surgeon can point to instruments and show where to cut.

“Remote experts can help people who have basic skills,” says Weibel. “With augmented reality goggles, we can present the expert as a hologram; we can even display their hands on the surgical field.”

A DESIGN-CENTERED APPROACH: With one foot in the Design Lab at UCSD, Weibel’s group naturally takes a design-centered approach, closely investigating problems before deciding whether technology can be part of their solutions.

“Research starts by observing the needs and stepping back to figure out the real problem,” says Weibel.

Sometimes, their findings surprise them. One of Weibel’s students spent a year in an ICU to learn about issues firsthand. What the ICU needed most, she realized, wasn’t augmented reality, but better family support. In the ICU, clinicians are constantly checking on patients while families watch in confusion; often they don’t know who the clinicians are or what they’re doing.

“Imagine having a phone that can simply tell families who entered the room and what kinds of questions they can answer,” Weibel wonders.

The lab prides itself on its unique blend of design-centered thinking and technological wizardry. In other projects, they are using video cameras to detect early symptoms of stroke and mining tweets to profile high-risk neighborhoods for HIV.

“We’re looking at all sorts of problems in health and healthcare,” says Weibel, “and we’re just trying to help wherever we can.”

The Search for Lifesaving Biomarkers

► **FOR PATIENTS BATTLING COLOR CANCER**, the years after surgery present a critical time for their recovery. Not only are these years especially prone to disease; in addition, a difficult decision must be made whether to undergo chemotherapy. Now these patients may have a new weapon in their arsenal.

The help comes from a team of researchers led by **Debashis Sahoo**, a professor in CSE and Pediatrics. The team discovered a so-called biomarker, or distinctive molecular feature, that identified colon cancer patients most likely to remain disease-free up to five years after surgery. Called CDX2, the protein biomarker also helped to identify stage II colon cancer patients who are most likely to benefit from chemotherapy after surgery. The retrospective study was published in 2016 by the *New England Journal of Medicine*.

“Because previous studies did not take into account differences between colon cancers with and without CDX2, doctors have long struggled to identify which stage II colon cancer patients might benefit from adjuvant chemotherapy,” said Sahoo. “But what we’ve now found is that some of these patients might benefit from chemotherapy, and we now have a biomarker to tell the difference, potentially saving many lives and reducing toxicity from unnecessary treatment.”

Model of differentiation in human colon crypt.



Debashis Sahoo

works in systems biology and bioinformatics. His recent work has used formal verification approaches to develop computational models of human normal and cancer tissues.



Sahoo led the study alongside co-first author **Piero Dalerba, MD**, of Columbia University, and senior author **Michael Clarke, MD**, of Stanford University.

The study took advantage of an earlier approach that Sahoo developed to identify differences in gene expression patterns. Sahoo pioneered this method to find genes involved in stem cell differentiation — the process by which stem cells specialize into specific cell types in an organ, such as the colon. “Dr. Sahoo’s bioinformatics approach is extraordinarily powerful,” says Dalerba. “We used it to search for biomarkers that could help us identify which colon tumors were likely to contain high numbers of stem-like cells.”

Dalerba and Sahoo discovered that when the gene CDX2 is “off,” another molecular marker of stem-like cells in colon tumors, called ALCAM, is always “on.” With this in mind, the team analyzed a database of cancer gene expression from more than 2,000 patients with known treatment courses and outcomes. They also used the database to study the association between CDX2 status and patient outcomes.

By examining data on 466 patients with colon cancer, the team discovered that CDX2-negative tumors were associated with poorer prognosis. They also found that treating CDX2-negative Stage II colon cancer patients with chemotherapy after surgery could improve their survival.

More work remains to be done. “While promising, this study was retrospective, meaning we looked back at existing patient data,” Sahoo says. “Before they can be applied to clinical practice, these results need to be confirmed by prospective, randomized clinical trials.” Nevertheless, the results mark a hopeful step toward better prognoses for this deadly disease.

Tracking the Skin’s Microbiome

► **TO UNDERSTAND HOW** microbes spread disease, a group of multidisciplinary researchers, led by **CSE Professor Nuno Bandeira**, set out to map the location of molecules on human skin. The skin belonged to two healthy adult volunteers, one male and one female, who had not bathed, shampooed or moisturized for three days. The researchers swabbed 400 different body sites on the volunteers and used mass spectrometry (MS) to determine the molecular and chemical composition of the samples. They also developed custom software to construct 3D models (see figure) of the data for each sampling spot.

Despite the three-day moratorium on personal hygiene products, the most abundant molecular features in the skin swabs still came from hygiene and beauty products, such as sunscreen.

The study was made possible by an infrastructure, called Global Natural Products Social Molecular Networking (GNPS), that Bandeira’s team developed. GNPS is an open-access knowledge base for community-wide organization and sharing of MS data.

“The algorithms and computational infrastructure in GNPS makes it possible to explore and navigate data from thousands of samples (each spot on each person’s skin is a separate sample), thereby enabling the whole community to add value and meaning to the data over time,” says Bandeira, who has a joint appointment in the Skaggs School of Pharmacy and Pharmaceutical Sciences. “Without this,

studies like this one would be hard, if not impossible, to conduct.”

GNPS includes a crowdsourced, wiki-like knowledgebase that accepts submissions and revisions of annotated MS data from any scientist. The submissions are immediately and freely available to the entire community.

Bandeira’s team was primarily responsible for developing the cyberinfrastructure and software tools to support GNPS, whereas his colleagues at the Skaggs School of Pharmacy and Pharmaceutical Sciences contributed the data and annotations.

The ultimate goal of this research is to understand the skin’s role in human health and disease. The maps generated by this study, combined with the GNPS infrastructure, provide a baseline for future work on the interplay between our personal hygiene routines, the molecules that make up our skin, and the microbes in our environment.



Nuno Bandeira is interested in algorithmic and statistical problems in proteomics.



First-ever maps of molecular diversity of male (left) and female (right) human skin.

Enlightened Reality

► **SOMETIMES IT TAKES AN EXPERT PHOTOGRAPHER** to light a difficult room or capture the beauty of a San Diego sunset. But that may change soon due to a series of breakthroughs by **CSE Professor Manmohan Chandraker**, who works at the intersection of computer vision and graphics.

“An image is the outcome of a complex interaction between the geometry of a scene, its constituent objects, and the illumination under which they are observed,” explains Chandraker. “At its core, our work is trying to unravel this process.” Chandraker wants to turn our camera phones into full-fledged devices for augmented reality. He sees a future in which any image or video can be photorealistically enhanced by non-expert users. It’s an ambitious goal.

To edit an image in a photorealistic manner requires estimating all the factors that affect its composition. Fundamentally, this is an ill-posed problem because different combinations of geometry, lighting, and materials can result in the same image.

Chandraker’s team is making progress by training deep neural networks on large datasets of annotated images. The work relies on simulations where computer graphics is used to render realistic images of scenes with complex material and lighting properties. To do this at scale, the team developed a new GPU-accelerated renderer that was an order of magnitude faster than the previous state of the art.



Manmohan Chandraker works on computer vision — in particular, 3D reconstruction, scene understanding and graphics-based vision, with applications in autonomous driving, robotics and augmented reality.

Faster rendering, however, wasn’t their only challenge. Despite the sophistication of computer graphics, a gap remains between computer-generated and real images. To bridge this gap, Chandraker’s team has been incorporating prior knowledge of the physical world into the design of deep networks. Specific layers in the networks are tasked with various sub-problems, such as rendering an image under new lighting conditions or modeling the global illumination when light reflects from one part of a scene to another.

Unlike previous methods, Chandraker’s networks learn to capture subtle effects of the real world — for example, how an object appears differently beside a dark corner or a bright window, or how its color bleeds into that of a painted wall.

“The human eye is sensitive to all these details,” says Chandraker, who speaks like a photographer looking for the perfect shot. “That’s why removing these artifacts is so important. For a true augmented reality experience, everything has to be just right.”

Original →



Augmented →



The Fast and The Furious

► **COMPUTER-GENERATED IMAGERY (CGI)** has come a long way since the days of Tron. But that progress has also raised expectations among today’s moviegoers, who are quick to notice when CGI falls short of its promise.

Meeting those expectations — and solving the technical problems that stand in the way — is the job of **CSE Professor Ravi Ramamoorthi**, who directs the Center for Visual Computing at UC San Diego. In fact, moviegoers have already seen Ramamoorthi’s work on the silver screen; it was used by Weta Digital for *War for the Planet of the Apes*, which was nominated for a visual effects Oscar.

A special area of interest for Ramamoorthi is the efficient rendering of fur. He, fellow **CSE Professor Henrik Wann Jensen**, and their students have developed a revolutionary approach to simulate the way that light bounces within an animal’s pelt. Their approach is dramatically faster and more accurate than the existing state of the art.

“Existing models were designed to create computer-generated hair,” says Ramamoorthi, who knows how to spot a fake fur. “But those models don’t take into account the central cylinder, or medulla, which is present in each fur fiber. The medulla in fur is much bigger than in human hair, and the appearance of fur arises from the particular way that light passes and scatters through that cylinder.”

Ramamoorthi’s approach pays special attention to the subsurface scattering of light — how it enters the surface of a translucent object at one point, scatters at various angles depending on the object’s material properties, and ultimately exits the object at a different point. The approach is especially good at modeling the way light bounces around fur fibers.

Subsurface scattering has been used previously in computer graphics and vision simulations, but this is the first time it’s been applied — with groundbreaking new results — to CGI furrries.

“Actors don’t have the luxury of bad hair days,” says Ramamoorthi, who sees many further applications for the technology, including video games and computer-generated special effects. “And when we’re done, neither will their CGI counterparts.”

Ramamoorthi is not lone in seeing this work as groundbreaking. His Ph.D student Lingqi Yan (now faculty at UC Santa Barbara) received the ACM SIGGRAPH Outstanding Doctoral Dissertation Award for his contributions to this effort.



Ravi Ramamoorthi researches the foundations of computer graphics and vision.



“Actors don’t have the luxury of bad hair days. And when we’re done, neither will their CGI counterparts.”





< I want to give programmers the tools to produce software faster and with fewer errors. />

<Programs that Write Programs Right/>

► **WHEN CSE CHAIR DEAN TULLSEN** introduced CSE Professor Nadia Polikarpova during a reception for new faculty, he joked that she was going to put software developers out of work. While this isn't quite true, Polikarpova does create programs that can generate other programs. But her goal is to help developers, not put them out of business. "Just like assembly lines in the auto industry made it possible to produce cars faster and more reliably, I want to give programmers the tools to produce software faster and with fewer errors," Polikarpova says.

Human developers make mistakes, and this has caused a number of major issues, from security breaches, to power outages, to malfunctions in cars. So Polikarpova's goal is to create the next generation of programming languages, which allow developers to directly state a system's high-level requirements. Then the programming language takes over and implements the code to make sure these requirements are met. This is known as "program synthesis."

Traditionally, building high-assurance software with provable reliability guarantees has been very expensive. But Polikarpova doesn't believe that there should be a tradeoff between program reliability and productivity and cost. She and colleagues have developed a number of programming languages to make it easier for programmers to build high-assurance software.

For example, Lifty helps app developers properly secure sensitive data, such as your credit card number, your location or the hidden parts of your social network profile. First, developers create the app, then they specify what the app's privacy policy should be. Finally, Lifty uses formal program verification technology to analyze the app's code and pinpoint all the places where sensitive data could involuntarily be leaked. The programming language then synthesizes the necessary access checks in the code in all those places. In addition, Lifty produces a mathematical proof that the resulting code will never leak sensitive data.

Polikarpova said she and colleagues are benefiting from advances in automated theorem proving and constraint solving algorithms. "We use computers' power to reason logically to generate programs that are guaranteed to be correct," she says. "But it's still impossible to quickly generate a complex program like Word from scratch." Developers, breathe easy... for now.



Program Verification Takes Flight

► **ERRORS IN CYBER-PHYSICAL SYSTEMS** can have disastrous consequences. In 1987, an integer overflow disabled safety checks on the Therac-25, an early machine for cancer treatments, subjecting patients to fatal overdoses of radiation. In 1996, a floating-point error caused the Ariane 5 rocket, valued at \$500 million, to explode in its maiden launch above French Guiana. More recently, software errors in Toyota vehicles have been linked to critical driving failures, leading to mass recalls and costly lawsuits.

To reduce the chances of such errors, CSE Professors Sorin Lerner and Ryan Kastner, along with MAE Professor Miroslav Krstic, are working to apply a formal technique, known as foundational verification, to cyber-physical systems. In foundational verification, a system is proved correct in full detail, using a proof assistant, before it is deployed in the real world. These formal proofs can guarantee previously unattainable levels of safety.

To test these ideas in practice, the researchers are implementing them on live drone — specifically, on flying quadcopters. They have developed a system, called VeriDrone, to ensure that drones operate safely and obey specific parameters in flight. For example, drones can be required to fly below the FAA-allowed altitude of 400 feet; they can also be forced to stay within a specific geographical area. One of VeriDrone's modules has been adopted by the popular Ardupilot open source project and installed on over 1 million vehicles worldwide.

The researchers are not only working with quadcopters in the lab, but also in the field. In particular, Kastner, with collaborators at San Diego Zoo Global, has developed a quadcopter-based system for tracking iguanas through the mangrove swamps of the Cayman Islands. The quadcopters are equipped with radio receivers and used to locate iguanas that are fitted with small radio collars. From the air, the quadcopters are able to locate iguanas much faster than ground-based observers with hand-held antennas. In this setting, there are several interesting properties to verify, such as whether the entire target geographic area has been flown over, or whether all the iguanas have been located with high confidence. These same properties can also be verified when multiple quadcopters are collaborating to cover an area.

Although the initial work focuses on quadcopters, the ideas in this research have far-reaching applications. Many kinds of systems — from automobiles and avionics to pacemakers and power grids — are likely to benefit one day from this approach.



Sorin Lerner develops programming language techniques for making software systems more reliable.



Ryan Kastner studies hardware acceleration, hardware security and remote sensing.

Architectures for Next-Level AI

► **CSE PROFESSOR HADI ESMAEILZADEH HAS BEEN FASCINATED BY AI** and machine learning ever since his days as an undergraduate. He and his colleagues have warned, however, that today's multicore and general-purpose processors will not be enough to power the next-generation applications of AI on our phones, laptops, and other devices.

"Two things have propelled AI and machine learning to the next level," Esmaeilzadeh explains. "One has been the advances in the algorithms, but the other has been the advances in the microarchitecture of processors. Algorithms in AI require massive amounts of computation, and increasing levels of performance will not be possible without changes in processor architectures."

Esmaeilzadeh first outlined this dilemma in 2011, and his work drew attention — not only from the world's leading researchers in computer architecture, but also *The New York Times*, which ran the headline "Progress Hits Snag: Tiny Chips Use Outsize Power."

For his leadership in the field, Esmaeilzadeh was named 2018's "Young Computer Architect" by the IEEE Technical Committee on Computer Architecture. He received the recognition at the 45th International Symposium on Computer Architecture.

Esmaeilzadeh also follows how companies are coping with these problems. Tech giants Microsoft and Google, for example, have turned to specialized chips and programmable accelerators for certain types of applications. But it can be difficult to maintain specialized hardware that keeps pace with the latest advances in AI and machine learning.

Esmaeilzadeh wants to make it easier. His approach is to develop what he calls algorithm-defined specialized computing stacks.

"We are going to the origins of these AI applications," he says, "and understanding the mathematical and theoretical foundations of their algorithms. When we understand that, we can design specialized programming languages, compilers, runtime systems, operating systems, and even microprocessor architectures."

"How do we bridge this gap between specialized hardware and these ever-evolving algorithms in AI?" Esmaeilzadeh asks. "This has been the challenge of my research."



Hadi Esmaeilzadeh works on computer architecture design for next generation algorithms, applications, and systems.

Revolutionary Computer Chip Design

► **CSE PROFESSOR ANDREW KAHNG HAS RECEIVED AN \$11.3 MILLION AWARD** from DARPA to develop electronic design automation tools for 24-hour, no-human-in-the-loop hardware layout generation. As director of the project, Kahng leads a team of researchers from multiple universities (UC San Diego, Brown, U Michigan, U Minnesota, and UT Dallas) and industrial partners (Qualcomm, Arm). **CSE Professors Lawrence Saul and C.K. Cheng** are co-PIs at UCSD.


Today, designing a new application-specific microchip can take 12 months or more.

"For the U.S. to be the vanguard of innovation we need to fully leverage semiconductor technology," Kahng says. "There's an incredible delta between what's possible with silicon versus what people are attempting. We're trying to narrow that gap."

Today, integrated circuit chips are designed by large teams of researchers with complex computer-aided tools. But the tools can be slow, and the problems of chip floorplanning, device placement, and routing remain formidable. Within teams, specialists are devoted to integrated circuits, systems-in-package, and printed circuit boards. Even with significant resources, the process can take upwards of a year.

Kahng wants to shorten that cycle to 24 hours, eventually without sacrificing power, performance, or quality. He also wants to put it in the reach of much smaller teams. To do so will require a new approach.

"If you want to finish in 24 hours, you have to free yourself from several long-standing mindsets," Kahng says. "We simply have to accept fragmentation of the design problem as a requirement, and then figure out how to perform high-quality design optimization within that constraint. Parallel and cloud-based approaches will be essential here."



Andrew Kahng's research focuses on the VLSI design-manufacturing interface, IC physical design, and combinatorial and global optimization. He pioneered methods that link designer intent with the manufacturing process to improve yield and reduce power.

FROM CSE 125 TO POKEMON GO

B Back in 2004, Ed Wu was a student in the first class to present the video game systems they created for CSE 125 in front of a camera. Fast-forward to today, and Wu is now a vice president at Niantic, the company that introduced augmented reality at worldwide scale to smartphone games with *Ingress* (2013), *Pokémon Go* (2016) and recently *Harry Potter: Wizards Unite* (2019).

Wu has had many pioneering roles at Niantic. He was a launch engineer for *Ingress* and the technical director of *Pokémon Go*, and now serves as the VP of platforms for all of Niantic's games. Wu regularly comes back to UCSD to present to CSE 125 students at the invitation of CSE Professor Geoffrey Voelker. He speaks about his work at Niantic and his experiences as an undergraduate.

"It's tough to describe how formative my experience here was," he says about his years as an undergraduate at UCSD. "I met my wife here. Some of my classmates became lifelong friends."

Wu was a Jacobs Scholar, as was his wife Simone. "UC San Diego was willing to invest in me and in other undergraduates," he says. "The university made it clear that they saw undergraduates as the lifeblood of the institution." The two later married in the Bear Courtyard in 2008.

As a student in John Muir College, Wu was a triple major in computer engineering, physics and political science. As an undergraduate, he worked on semiconductors in the research group of former UC President and UCSD physics Professor Robert Dynes. "He was an incredible inspiration," Wu says of Dynes.

After graduating from UCSD in 2004, Wu decided to pursue a Ph.D. in physics at Stanford. Wu focused on astrophysics and the cosmic background radiation that allowed researchers to determine the age of the universe. He then took a job with the RAND Corporation as an associate physical scientist, a position combining his interests in technology and political science. Meanwhile, his wife Simone, who is also an engineer, had landed a job at Google. She was inspired by her work there, and so Wu soon followed her. He worked at what would later become Niantic when the company spun off from Google in 2015.

Building video games teaches you about coding and serving data to hundreds of thousands of users. It also teaches you teamwork, Wu says. "Real software engineering happens in groups of half a dozen people or so," he says.



Ed Wu
is vice president
of Platform at
Niantic, and was
director for their
Pokémon Go
game.

“

Gaming is one of the manifestations of what computer science is capable of. It pushes you and challenges you on so many axes.”

—Ed Wu tells CSE 125 students

Niantic's mission is to convert people's relationship with the world so that they start interacting with one another rather than being isolated. *Pokémon Go* revolutionized augmented reality mobile games that incorporate user location as a fundamental aspect of the gaming experience. Users catch Pokémon, get supplies and battle in gyms on their smartphones, at any time and in any place from the United States, to France, to Australia. The game has been downloaded over 1 billion times worldwide and has nearly 150 million active users. Niantic's recent release of *Harry Potter: Wizards Unite* continues to innovate in location-based AR games, overlaying the universe of Harry Potter on the physical world surrounding players.

Creating a shared, consistent world for hundreds of millions of users is no easy feat. But Wu and his team had already honed their skills on *Ingress*, Niantic's first game. At the core of both games is Google Earth, which was Niantic's original product, back then known as Keyhole.

Niantic's games are free-to-play because the company wants to make sure that players can enjoy their game for free. "If you respect the users, they will respect you back," Wu says. "It's possible to play the game and have a lot of fun without having to spend money."

When speaking in CSE 125, he also dispenses valuable professional advice. For undergraduates who want to pursue a job in the gaming industry, Wu advises getting strong computer science fundamentals, including graphics, cloud computing and networking at scale. "You also need to have a passion for coding, for making your own games, and for making your own apps." He also encourages everyone to always learn more. Recently Wu earned an MBA from UC Berkeley. "I discovered I have a lot to learn, and not just technology," he says. "I want to be a well-rounded individual."

"At UC San Diego, I put down the foundations for a lifetime of learning and growing," Wu says.

Early Research Scholars Program Expands Opportunity

► FOR YEARS, THE NATIONAL SCIENCE FOUNDATION has issued a challenge across all fields of science and engineering to create a more diverse workforce. In CSE, that effort begins by recognizing that prospective majors come to the department with widely varying backgrounds. Computer education is not part of the standard curriculum in K-12 schools, and there can also be cultural differences in exposure to computer science and research.

“Students can be turned off for many reasons,” says **CSE Professor Christine Alvarado**. “Often they look around and don’t see any other students who look like them. Other times they are intimidated by students with vastly more experience.”

Alvarado started the Early Research Scholars Program (ERSP) in 2014 to address these challenges. The program has the unique goal of engaging second-year CSE students in a team research experience. Students apply for academic yearlong apprenticeships with CSE faculty mentors to design and complete independent research projects. To date, these mentors have supervised 185 students in 48 projects.




Christine Alvarado is developing curricula that make CS more engaging.

In addition to technical mentoring from their research advisers, students also receive general mentoring and support from Alvarado and her graduate student assistant. “The program has created a huge uptick in the number of students who can do research,” says Alvarado, “and the faculty mentors in ERSP have been impressed with the quality of the students and their accomplishments.”

Adrian Mendoza, a first-generation student whose parents are Mexican immigrants, worked on an application of face recognition during his ERSP year. He credits ERSP for introducing him to research and sparking his interest in graduate school. “ERSP opened new doors for me,” he says. “I was able to apply what I learned in class to real-world problems.” Mendoza was the 2018 recipient of the department’s Alan Turing Memorial Scholarship.

In 2018, the program received a \$2 million, five-year grant from the National Science Foundation to support its expansion to seven new universities, beginning with UC Santa Barbara, Stanford University and the University of Illinois at Chicago.

Alvarado looks forward to replicating the program’s success on a larger scale. “Early research can build students’ confidence and help them develop their identities as computer scientists,” she says. “And for some students, research becomes a lifelong passion. It worked here, and it can work other places, too.”



← Pictured left to right: Professors Leo Porter, Mia Minnes and Christine Alvarado.

One of the most popular features is a video series called *Concept Challenges*. Students answer questions on a topic, watch UCSD students discuss the question, and then listen to an instructor’s explanation. Afterwards, they get to try the original questions again.

When I Struggled, another video series, features UCSD faculty and students talking about the challenges they faced while learning challenging concepts.

“The goal was to help folks, primarily from underrepresented minority backgrounds, see everyone struggles learning this content at times,” says Porter.

Minnes notes she and her colleagues use these MOOC resources in their brick and mortar classrooms, including video interviews and panels. Prospective and transfer students also look up Coursera classes led by UCSD faculty. “I’ve heard a lot of appreciation from our current students for the online experiences they had before they got here,” says Minnes.

Porter says teaching the online courses is extremely rewarding in part because they reach underserved students. “Many are from countries that lack access to similar courses,” says Porter, “and many are full-time workers trying to learn a new skill to change careers.”

One student comment in particular drives this point home. “I really appreciate all the hard work you put into this course and all the other courses in the specialization. You are helping people all over the world in making a better life for themselves.”

Serving Up a Five-course MOOC

► TO EXPAND UCSD’S ACADEMIC REACH, and offer resources that transcend geographic boundaries, three CSE professors have created a five-course specialization on Coursera, the massive online open course (MOOC) platform. Called Object Oriented Java Programming: Data Structures and Beyond, the specialization helps participants build their interactive software development skills and has enrolled more than 295,000 students.

CSE Professors **Christine Alvarado**, **Mia Minnes** and **Leo Porter** created the curriculum in response to an RFP from Coursera, with Google contributing ideas for real-world projects and the involvement of its engineers as guest lecturers.

The courses cover intermediate Java programming and mastering software engineering interviews. All have been ranked 4.6 or higher (out of 5) by Coursera students.

Minnes says building the online specialization was a great opportunity for the instructors “to break down the barriers between courses and curricula and prerequisites and just say, ‘Let’s think about a really interesting application and walk someone who is curious about that application through building a project they find meaningful and lets them explore the space.’”



Leo Porter works in computer architecture and computer science education research.



Mia Minnes has expertise in computation and foundational mathematics.


Neha Chachra

Combating the Economic Pull of Cybercrime

► **COMPUTER SCIENCE WAS AN EASY CHOICE** for Neha Chachra when she began her studies, she says. “Growing up, I had an aptitude for math, sciences and logic.” But security wasn’t on her radar until a summer internship at Google in Mountain View exposed her to kinds of problems that would define her career. “That was really my first introduction to various attacks and how they subvert well-intentioned code, and I thought it was beautifully complex,” she recalls. Today, Chachra works as an engineering manager at Facebook, where she has been employed since 2016, leading a team that protects Facebook against abusive links, protecting the platform from harmful URLs that expose users to phishing, spamming and malware. The work is a natural progression of her research and interests at UCSD, where she received her Ph.D. in 2015.



Coming to UCSD was one of the best decisions I ever made.”



Neha Chachra works as an engineering manager at Facebook and serves on the UCSD CSE Alumni Board.

She was drawn to UCSD’s research program as a way to explore her combined interest in economics and computer security. “Fortunately for me, Professors Stefan Savage and Geoffrey Voelker had just started looking into the economics of spam,” says Chachra. “Working with them was one of the best decisions I ever made, as they turned out to be the most brilliant, yet kindest, advisers one could hope for.”

At UCSD, Chachra delved into research centered on understanding attacker ecosystems — namely, how people profit from cybercrime and spam. Her thesis “*Understanding URL Abuse for Profit*” looked at how attackers mounted various kinds of affiliate marketing fraud and the effectiveness of different interventions — both existing and proposed — in undermining their profitability.

“Today,” she says, “the sum total of all my research and that of my peers [at UCSD] continues to guide me as I lead the engineering team at Facebook, working on detection and the mitigation of abusive URLs.”

The work requires Chachra to consider both technical solutions as well as a little psychological jiu-jitsu. “Given two solutions, it’s critical that I think about how attackers might respond to either solution, and pick the more resilient option,” she says. “In an odd way, to be successful in this field, one frequently learns to think like an attacker.”

“The most interesting aspect of all cybercrime is that there is a real human being on the other side making rational choices to maximize profit as you or I would in any business,” Chachra adds. “Except, unfortunately, these criminal activities hurt the people using the Internet. The difficulty of solving problems in the presence of a human adversary, coupled with the importance of this work in protecting people, is what truly keeps me working on it.”

Sarah Meiklejohn Ph.D. 2014



Meiklejohn has broad research interests in computer security and cryptography, and works on topics such as anonymity in

cryptocurrencies, privacy-enhancing technologies and bringing transparency to shared systems. In addition to her research on cryptographic primitives and zero-knowledge proofs in particular, her dissertation work introduced new techniques for tracking the flow of money within the Bitcoin cryptocurrency. These techniques have since been adopted in both the public and private sectors, and her dissertation won the Chancellor’s Dissertation Medal. She is now an Associate Professor in the Department of Computer Science at University College London.

“The best part of having done my Ph.D. at UC San Diego was the open and encouraging atmosphere, which I have to admit I now feel really spoiled by. I was continuously provided with opportunities to learn, grow and get involved in the many exciting activities going on in the department.”

Zachary Tatlock Ph.D. 2014



Tatlock works on improving software reliability by developing tools that help programmers ensure their code is safe and

accurate. During his graduate training at UCSD, Tatlock and colleagues developed new techniques to mathematically prove critical software like Web browsers are secure; e.g., that they won’t leak secret cookies from one website’s tab to another site. Building on this experience, Tatlock has gone on to establish several “verification firsts” with colleagues at UCSD and the University of Washington, where he is now an Assistant Professor in the Paul G. Allen School of Computer Science and Engineering.

“Studying at UC San Diego absolutely changed my life. I got the opportunity to

learn from, hack alongside and develop deep friendships with the best researchers in my area. It all led to my dream job where I have the opportunity to pay it forward by spreading UC San Diego culture to help create that experience for others.”

Zachary Lipton Ph.D. 2017



Lipton’s work addresses core technical challenges and real-world applications of machine learning.

He collaborates closely with Amazon, where concurrent with his final year of his Ph.D., helped to grow the AI team at Amazon Web Services (AWS) from a five-scientist into a large applied research organization, and contributed to the Apache MXNet deep learning framework. Lipton is an Assistant Professor at Carnegie Mellon University in the Tepper School of Business, with appointments in the Machine Learning Department and the Heinz College of Information Systems and Public Policy.

“UC San Diego is the perfect place to pursue a Ph.D. in machine learning. The school boasts machine learning faculty as strong as at any department in the nation while offering the laid-back atmosphere and idyllic weather that one naturally associates with San Diego. Coming in as a musician, I found that the Ph.D. program offered me freedom to grow without looking over my shoulder, access to inspiring professors, and the freedom to take academic risks, build my own research agenda and chase problems across interdisciplinary boundaries.”

Sarah Guthals Ph.D. 2014



Guthals focuses on making technology more accessible to novice audiences by building engaging learning experiences,

founding companies focused on novice coders and writing technical books for the novice audience. Her dissertation

focused on the importance of teaching about the culture of computer science, not just teaching them facts and skills. By engaging novices (as young as 5) in what it means to be a computer scientist, she has helped improve accessibility. She is now a Senior Program Manager in the Academic Ecosystems team at Microsoft, focusing on bringing machine learning and data science skills to K-12 students.

“First as an undergraduate back in 2006, then as a graduate student and finally as a lecturer, UC San Diego, and specifically the Computer Science Department, has always been a place of inspiration, camaraderie and happiness for me in my life because of the incredible people. I continue to be inspired by the dedication this department has to the students and vice versa.”

Jason Oberg Ph.D. 2017



Oberg’s dissertation work was an intersection between computer architecture, hardware design and automation

specifically focused on detecting and preventing vulnerabilities in hardware. He is a co-founder and Chief Executive Officer of Tortuga Logic. Tortuga Logic is a cybersecurity company building products to detect and prevent hardware security vulnerabilities. He started Tortuga Logic during his final years at UCSD using valuable resources, including the National Science Foundation’s I-Corps program and UCSD’s Institute for the Global Entrepreneur.

“The CSE department fosters deep amounts of collaboration across every discipline. Beyond just my doctoral work, I learned an enormous amount from my friends and colleagues during my tenure in the department, much of which has been tremendously valuable as I continue to work with people across many technical disciplines. Beyond just the department, UC San Diego is a hotbed for innovation with access for students to start companies, find early investment and pursue applications of their research.”

CSE FACULTY

SYSTEMS & NETWORKING

ARTIFICIAL INTELLIGENCE



TAYLOR BERG-KIRKPATRICK



KAMALIKA CHAUDHURI



GARY COTTRELL



SANJOY DASGUPTA



YOAV FREUND

COMPUTATIONAL BIOLOGY



VINEET BAFNA



NUNO BANDEIRA



MELISSA GYMREK



PAVEL PEVZNER



DEBASHIS SAHOO

HUMAN-COMPUTER INTERACTION



WILLIAM GRISWOLD



SCOTT KLEMMER



LARRY SMARR



NADIR WEIBEL

ARCHITECTURE



HADI ESMAEILZADEH



PAT PANNUTO



STEVEN SWANSON



DEAN TULLSEN



JISHEN ZHAO

CS EDUCATION



CHRISTINE ALVARADO



JUSTIN ELDRIDGE



AARON FRAENKEL



MIA MINNES

SECURITY & CRYPTOGRAPHY



MIHIR BELLARE



NADIA HENINGER



DANIELE MICCIANCIO



STEFAN SAVAGE



DEIAN STEFAN

GRAPHICS & VISION



MANMOHAN CHANDRAKER



ALBERT CHERN



DAVID KRIEGMAN

EMBEDDED SYSTEMS



RAJESH GUPTA



RYAN KASTNER



PRADEEP K. KHOSLA



TAJANA ŠIMUNIĆ ROSING



JOSEPH PASQUALE



GEORGE PORTER



AARON SCHULMAN



ALEX C. SNOEREN



GEOFFREY M. VOELKER



YIYING ZHANG



YUANYUAN (YY) ZHOU



SICUN GAO



JULIAN MCAULEY



NDAPA NAKASHOLE



ALON ORLITSKY



LAWRENCE SAUL



SHANKAR SUBRAMANIAM

ROBOTICS



HENRIK I. CHRISTENSEN



LAUREL RIEK

PROGRAMMING SYSTEMS



RANJIT JHALA



SORIN LERNER



NADIA POLIKAROVA

THEORY



RUSSELL IMPAGLIAZZO



DANIEL KANE



SHACHAR LOVETT



RAMAMOHAN PATURI



ALEXANDER VARDY



NIEMA MOSHIRI



JOE POLITZ



LEO PORTER



GERALD RAJ

DATABASES



ALIN DEUTSCH



ARUN KUMAR



YANNIS PAKONSTANTINO



VICTOR VIANU



RAVI RAMAMOORTHI



HAO SU

VLSI/CAD



C.K. CHENG



ANDREW KAHNG



ALEX ORAILOGLU