e&s
Engineering & Science
California Institute of Technology
Pasadena, California 91125

Engineering & Science is printed on paper from responsible sources in accordance with the Forest Stewardship Council's sustainability standards.

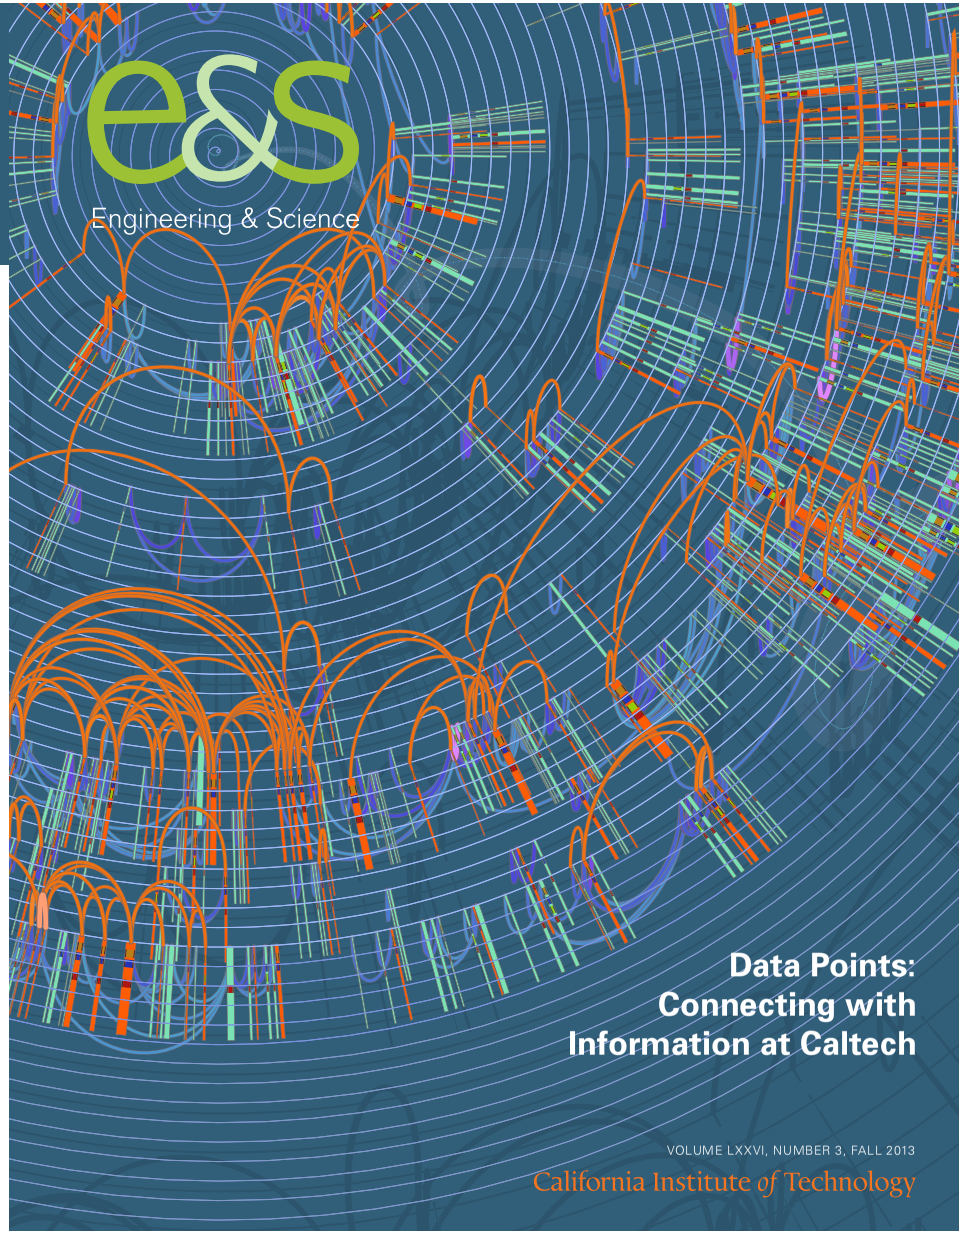

Data Points: Connecting with Information at Caltech

VOLUME LXXVI, NUMBER 3, FALL 2013
California Institute of Technology

e&s
Engineering & Science
VOLUME LXXVI, NUMBER 3, FALL 2013

2 Caltech on Twitter

3 From the President

4 Random Walk
Stalagmite Climate Records ▪ The Art of Science ▪ Ditch Day ▪ And More

9 Available Now on Caltech.edu

10 Seeing Is Knowing
BY KATIE NEITH
Caltech scientists and engineers take imaging to the next level, pushing the boundaries of how we see the physical environment around us and how we visualize the building blocks of that physical world.

16 What's the Big Deal about Big Data?
BY MARCUS Y. WOO
A look at the impact big data is having on science and how Caltech computer scientists and engineers are creating the techniques and infrastructure that will be needed for us to navigate our data-intensive future.

24 Navigating the Brain's Mysteries
BY KIMM FESENMAIER
Caltech researchers (and the White House) say it's time to piece together a dynamic map of the brain—one that shows its complex trafficking across trillions of neuronal connections. Addressing this grand challenge could just be the technological moon shot of a generation.

28 Evolving Education
BY ANDREW ALLAN
Training the scientists and engineers of the future starts with teaching for the future.

32 Alumni Impact: Game Theory, Austen Style
BY THE CALTECH ALUMNI ASSOCIATION

35 In Memoriam
Francis H. Clauser
Donald Coles

36 Endnotes
Summing up the Caltech experience.

ON THE COVER:
Computational scientist Santiago Lombeyda created this looping, swirling piece of art (and the more conventional diagram seen on these pages) from the masses of data on gene interactions coming out of the lab of biologist Barbara Wold, whose work is described on page 21. As part of Caltech's Center for Advanced Computing Research, Lombeyda serves as both data visualization guru and artist, creating computer models, simulations, and more for scientists and engineers across all disciplines on campus.

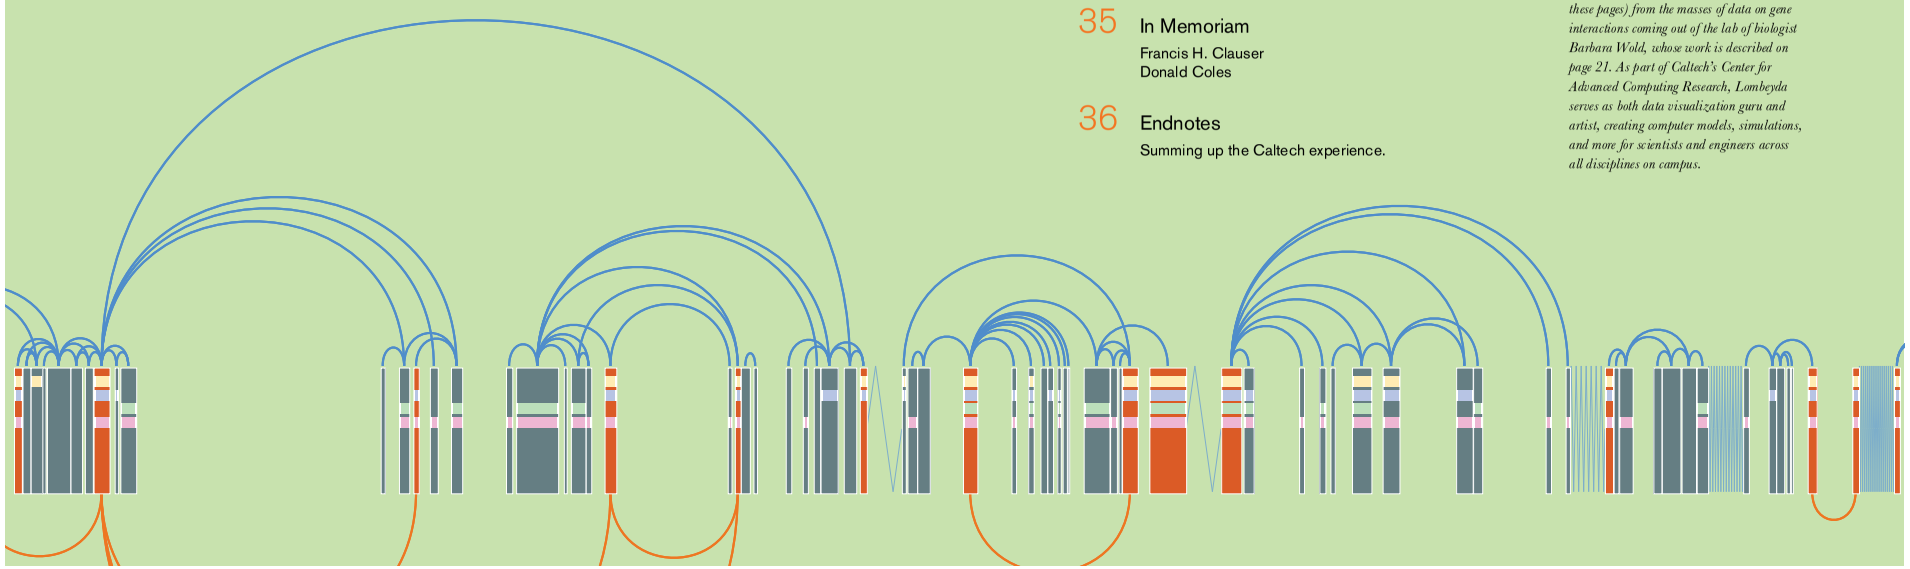

Caltech on Twitter

Follow us, retweet us, and let us know you're talking about us by including @Caltech in your tweets.

@Jordan2Mars: You know you're at a Caltech party when people pause to see the International Space Station. #1yearonMars pic.twitter.com/1uHQ7r9O4b

@loveplanets: Just saw a dude whose shirt said, "Flirt harder. I'm a physicist." I'm not sure if I feel attracted or repulsed. #Caltech

@NBCLA: An ancient ocean may have covered 1/3 of Mars' northern hemisphere, @Caltech scientists find. http://4.nbcla.com/1aRvwbt pic.twitter.com/a2GVtTWcvj

@CFCamerer: Nick Christakis in NYT advocates more lab experiments to teach economics+. @ Caltech we have done so for decades! http://nyti.ms/12Tz4zP

@JadenGeller: I'm usually not a sunglasses guy, but I really like wearing these @Caltech Fund glasses they sent me.

@Nighthawk_Black: Exciting to think the location of the first earth analog we're going to find may already be sitting in this archive: http://exoplanetarchive.ipac.caltech.edu/index.html

@konradyandson: This is what you call innovation! Check out this house by @sciarc and @Caltech http://meetdale.com/

@HannaStorlie: Feast your eyes on the @Caltech E&S summer publication. #DreamJob

@RileyEDixon: Fell in love with Caltech's environment but it's EXTREME! An amazing school to say the least.

Tweets may have been edited for spelling and grammar.

Engineering & Science
VOLUME LXXVI, NUMBER 3, FALL 2013

EDITOR
Lori Oliwenstein

MANAGING EDITOR
Doreese Norman

ASSOCIATE EDITOR
Katie Neith

CONTRIBUTING WRITERS
Andrew Allan, Kimm Fesenmaier, Katie Neith, Ben Tomlin, Marcus Y. Woo

COPY EDITORS
Michael Farquhar, Sharon Kaplan

ART DIRECTOR
Jenny K. Somerville

STAFF ARTIST
Lance Hayashida

DESIGN AND LAYOUT
Keith & Co.

BUSINESS MANAGER
Rosanne Lombardi

THE CALIFORNIA INSTITUTE OF TECHNOLOGY
Brian K. Lee, *Vice President for Development and Institute Relations*

Kristen Brown, *Assistant Vice President for Marketing and Communications*

Binti Harvey, *Director of Institute Marketing*

THE CALTECH ALUMNI ASSOCIATION
Heather Dean, *President*

Read *E&S* on the go at EandS.caltech.edu

Contact *E&S* at lorio@caltech.edu

Engineering & Science (ISSN 0013-7812) is published quarterly at the California Institute of Technology, 1200 East California Boulevard, Pasadena, CA 91125. Annual subscription $20.00 domestic, $30.00 foreign air mail; single copies $5.00. Send subscriptions to *Engineering & Science*, Caltech Mail Code 5-32, Pasadena, CA 91125. Third-class postage paid at Pasadena, CA. Published by Caltech and the Alumni Association. Telephone: 626-395-8371.

Printed by Navigator Cross-Media, Inc., Glendale, California.

PICTURE CREDITS
Inside back cover, TOC, 20 (bottom) – illustration by Santiago V. Lombeyda, Center for Advanced Computing Research, based on data courtesy of the Wold Lab and tools developed by Lombeyda and Katherine Fisher-Aylor of the Wold Lab; 3, 4–5, 7 (top left and right bottom two graphics), 18, 19 (bottom), 20 (left), 22 (top), 24 – Lance Hayashida; 6 – Jenny K. Somerville; 7 (top right), 10, 12, 15 – Shutterstock; 7 (bottom) – micrograph image by Mary Kennedy, fiber collage created by Liddy Hubbell; 8 (left) – Xun Wendy Gu; 8 (right) – Syria Lejau; 9 – B. Saxton & A. Angelich/NRAO/AUI/NSF/ALMA (ESO/NAOJ/NRAO); 11–12 – Everett Kane; 12–13 – Ahmed Zewail and Ulrich Lorenz; 13–14 – Kaushik Sengupta; 16 – visualization by Santiago V. Lombeyda, simulation by Bo Li, Michael Ortiz, and Caltech's PSAAP Center; 19 (top) – L. Shyamal via Wikimedia Commons; 22 (bottom) – Caltech Athletics; 23 – image courtesy of the CRTS survey team; 26 – iStockphoto; 28–31 – Martin Springborg; 32 – Dave Lauridsen; 35 (top) – Don Ivers, courtesy of the Caltech Archives; 35 (bottom) – EAS Communications Office; 36 – Richard Hsu

Special thanks to Maria M. Nastasescu

Connecting with Information

"What do the data tell us?"

This is one of the most common questions you will hear at Caltech, reflecting the fact that observations of nature and of experiments, and the analysis of these data, are the primary bases of modern science and engineering. But what if you have a set of observations for which there is no known method for extracting, recording, mapping, or sharing useful information? Or what if the sheer volume of information makes it difficult to discern patterns or to distinguish a true signal from noise?

These issues in science and engineering are becoming even more essential given the volumes of data now available and the small signals generated during critical observations and experiments. Thus, an aspect of modern research is the development of novel ways to look at and learn from new forms and volumes of information. Caltech students and faculty are working on new methods for large-scale data collection, transfer, and analysis (see page 16); for creating and utilizing visualization techniques that produce images bordering on fine art (see page 10); and for mapping the connections and interactions that enable our ultimate information gatherer—the human brain (see page 24).

Much of our work in training the next generation of scientists and engineers…is about teaching them how to mine and interpret data in such a way that it becomes useful and meaningful.

Much of our work in training the next generation of scientists and engineers—in training critical thinkers of all kinds—is about teaching them how to mine and interpret data in such a way that it becomes useful and meaningful. In this way, education at Caltech is as much about how to use and manage information as it is about the principles and facts underlying modern science and engineering. And as we find new and better ways to communicate about data (and use data to communicate), we are seeing its impact extend from the lab into the classroom (see page 28).

Our students are being shaped by the information they are learning in Caltech's classrooms and by the data they are gathering in its labs. This issue of *E&S* provides a glimpse into what may well be looked back on as a period of transition in science, engineering, and education.

Sincerely,

Edward Stolper

Random Walk

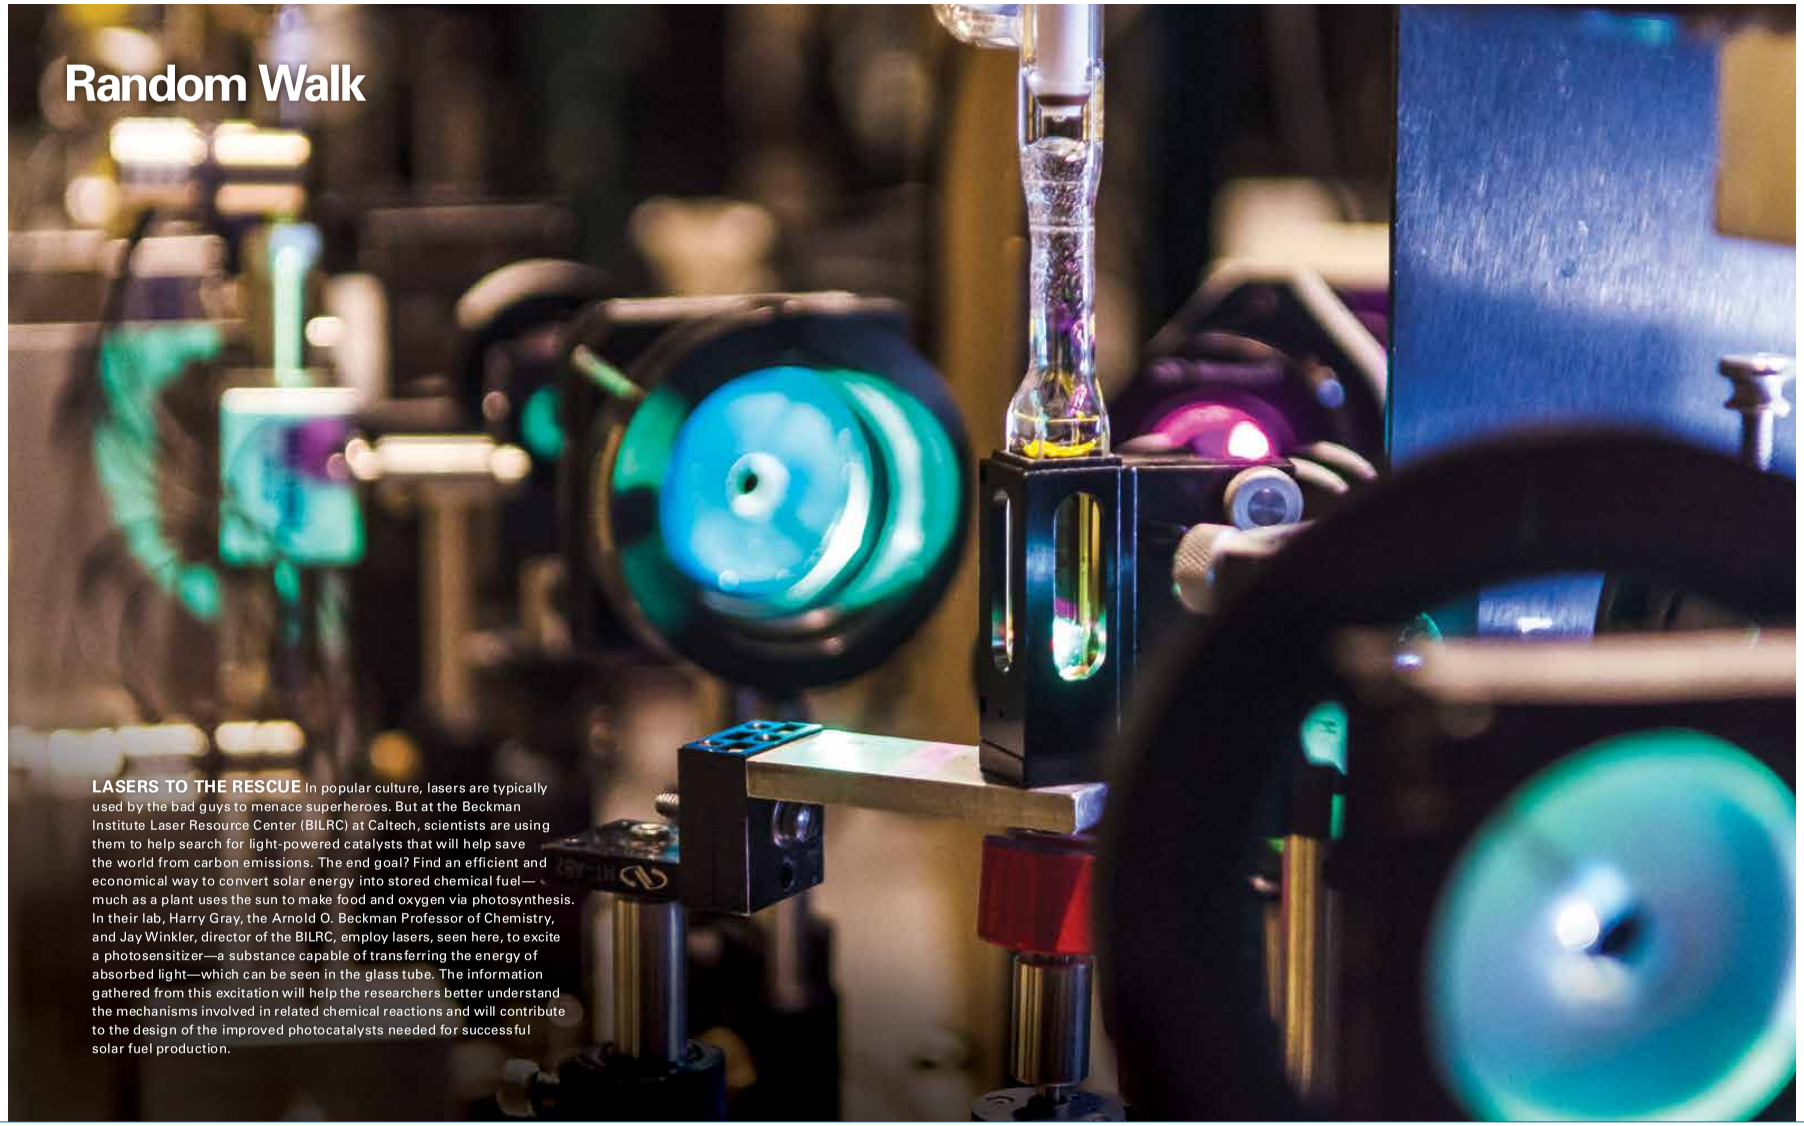

LASERS TO THE RESCUE In popular culture, lasers are typically used by the bad guys to menace superheroes. But at the Beckman Institute Laser Resource Center (BILRC) at Caltech, scientists are using them to help search for light-powered catalysts that will help save the world from carbon emissions. The end goal? Find an efficient and economical way to convert solar energy into stored chemical fuel—much as a plant uses the sun to make food and oxygen via photosynthesis. In their lab, Harry Gray, the Arnold O. Beckman Professor of Chemistry, and Jay Winkler, director of the BILRC, employ lasers, seen here, to excite a photosensitizer—a substance capable of transferring the energy of absorbed light—which can be seen in the glass tube. The information gathered from this excitation will help the researchers better understand the mechanisms involved in related chemical reactions and will contribute to the design of the improved photocatalysts needed for successful solar fuel production.

The HTML of Systems Biology

An international team of scientists recently published the most comprehensive virtual picture of human metabolism ever constructed. The computational model, called Recon 2, stitches together several previous models of metabolism along with many new research findings, and it includes more than 7,400 biochemical reactions and 5,000 metabolites. The hope is that the new tool will help scientists identify previously unknown causes of diseases, make predictions of drug efficacy, and even enhance personalized medicine.

But how did they combine all that data, considering that each piece was encoded by different software? The massive reconstruction was made possible by the work of Mike Hucka, a Caltech staff scientist in computing and mathematical sciences. About 13 years ago, Hucka—working in the lab of John Doyle, the John G Braun Professor of Control and Dynamical Systems, Electrical Engineering, and Bioengineering—first developed a machine-readable common format for computational models called Systems Biology Markup Language (SBML), which is a bit like HTML.

"It's hard to send an entire database to somebody else's program," Hucka says. "You need a way to write it out—an independent common exchange language—and SBML provides that."

Over the years, SBML has enabled the creation of nearly 1,000 large-scale models, not only of metabolic networks but of everything from neural processes and cellular signaling to a zombie apocalypse. (OK, a model of the way infectious diseases spread.)

"This way of writing down a model in a formal form allows people to do better science," Hucka says. "It makes models more precise, and it enables many people to collaborate on the same problem, which is something that people have to do more and more, especially in biology."—*KF*

Cities are responsible for **70%** of fossil-fuel **CO_2 emissions**.

To address this problem, the newly created Megacities Carbon Project will develop and test methods for monitoring the greenhouse gas emissions of not only cities—such as Los Angeles and Paris, France—but also their power plants. The project includes scientists from Caltech's Keck Institute for Space Studies (KISS), NASA's Jet Propulsion Laboratory, and other collaborating institutions. For more information, visit the project's website at megacities.jpl.nasa.gov/portal.

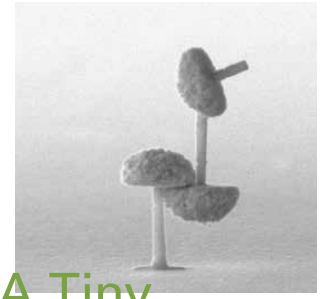

A Tiny Circus Act

These nano-mushrooms are made of nanocrystalline platinum, a precious metal that can be used in everything from electrodes and dentistry equipment to expensive jewelry. In this case, the platinum mushrooms are actually nanoscale testing specimens that help scientists understand the effects of sample-size reduction on the mechanical strength and likelihood of fracture of the material being tested. In this image, captured by Xun Wendy Gu, a graduate student in the lab of Julia Greer, professor of materials science and mechanics, the mushroom stalks are just 120 nanometers in width, with a single nanometer being equal to one-billionth of a meter. To test the platinum, diamond grips are used to pull at the mushroom heads. This sometimes separates the mushrooms from the material on which they are embedded so that the mushrooms stick to each other in visually interesting ways. Gu's image was so interesting, in fact, that it won first place in the 2013 Art of Science competition, a semiannual event that celebrates the interplay between art and science. This year's exhibit, featuring 35 images, was organized by Caltech's literary and visual arts magazine, *Totem*, and sponsored by Robert Gerson Metzner (BS '38).—*KN*

MIGHTY Stalagmites

For environmental scientists, ice cores and ocean sediments have long served as important indicators of past climate change, providing records of precipitation, temperature, atmospheric conditions, and more for particular periods in our planet's history. It was believed that for certain large-scale climate changes, these weather archives reflected conditions that applied to all of the earth, but new findings from stalagmites—large fingerlike formations that rise from the floors of limestone caves—imply that the planet's tropical regions may have a climate history all their own.

According to a study by Caltech geochemist Jess Adkins and colleagues at the Georgia Institute of Technology, climate records found in stalagmites gathered from Borneo indicate that the western tropical Pacific responded very differently than other regions of the globe to abrupt climate change—a finding that could help researchers better understand what might happen during future climate change events.

Stalagmites are formed when rainwater seeps into the ground, dissolving limestone rock that drips into the caves to grow the structures at a rate of roughly one centimeter every thousand years. Adkins worked to correlate different depths in stalagmite samples with different periods of time in much the same way scientists date trees by analyzing their rings; he used a type of radioactive dating technique called U-series dating, then merged data from four different stalagmites from caves across Borneo. Together, these analyses provided a record of precipitation trends in the western Pacific over the past 100,000 years.

What they found, says Adkins, is that "some historic climate changes, defined by records from ice cores and ocean sediments in the high latitudes, show up in Borneo, but others do not. This is surprising because most of the field thinks that the mechanisms behind events that happen at higher latitudes are similar, if not identical, to those in the tropics. Our findings question these notions." —*KN*

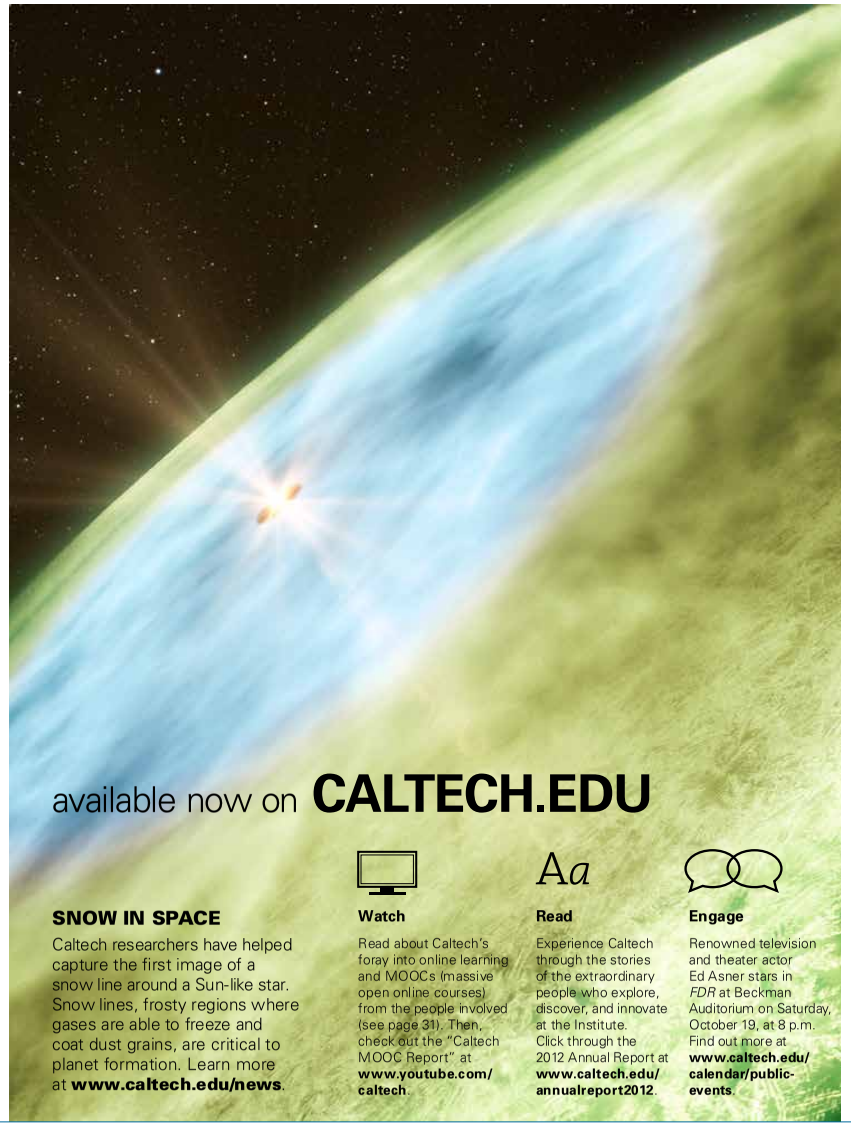

Seeing Is Knowing

Caltech scientists and engineers create new lines of sight with innovative imaging techniques

BY KATIE NEITH

Photographs add emotional context to current events, bring loved ones to mind with a single glance, and keep the past alive when our memories begin to fade. In science and engineering, still and moving images provide a similarly essential service, offering insight into—and sometimes even answers for—the biggest questions researchers are asking about how the world works.

At Caltech, a number of scientists and engineers are opening up new views into our environment and are exploring how we can best visualize the structures and processes that make up the chemical compounds and tiniest biological elements of the physical world.

"There is so much about cells that we don't understand," says structural biologist Grant Jensen. "But we're learning that many biological processes can be understood just by imaging the machinery working in a cell, doing its job. Just a picture of it reveals basically how it works."

Up Close and Biological

Jensen says that, growing up near Los Alamos National Lab in New Mexico, "I imagined I would become a physicist like everybody else in town." But soon after he started his undergraduate studies in physics at Brigham Young University, a research project led to his being captivated by the wonders of and potential advances to be made in structural biology. He never looked back.

"I learned that some of the highest-impact work waiting to be done in biology involved image processing, math, and three-dimensional reconstructions, which is all stuff I love," Jensen recalls. "As a postdoc, I saw that it was possible to take multiple images of a single object from different points of view and merge them into a 3-D reconstruction, and I realized that we could do this to cells. Immediately, I knew that this was what I would be doing for many years."

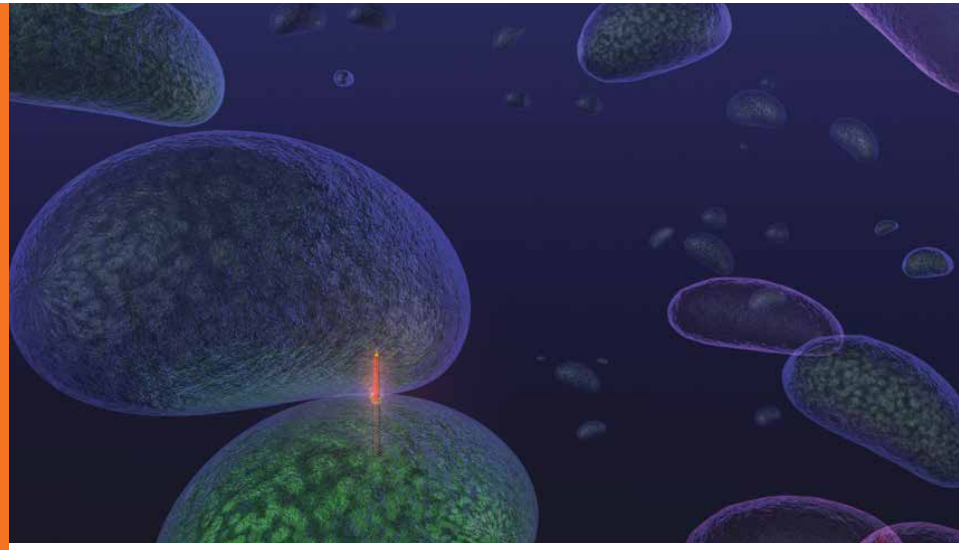

Jensen says his imaging research is guided by the words of an unlikely scientific muse, baseball's Yogi Berra, who once famously said: "Sometimes you can see a lot just by looking." To do that looking, his lab has become one of just a handful in the world to own and operate an electron cryomicroscope—a unique type of microscope that enables a novel imaging technique called electron cryotomography (ECT), which allows Jensen and his team to observe biological samples in a near-native state. Developed in Europe in the 1980s, the electron cryomicroscope produces a magnified image by illuminating samples—kept at cryogenic temperatures of below –150 degrees Celsius—with an electron beam. Unlike in traditional microscopy, for which samples must be fixed, embedded in plastic, sectioned, and stained, samples for ECT are frozen so quickly that they become almost immediately fixed within a layer of transparent, glass-like ice.

Once the sample is frozen, it is rotated around an axis while a specialized digital camera takes a series of high-resolution images; the information gathered by the camera allows the team to reconstruct the object in three dimensions and then analyze it in detail.

"Caltech is positioned to do this kind of work because a generous gift from the Gordon and Betty Moore Foundation allowed us to buy the world's very best electron cryomicroscope in 2002," Jensen says. "Only one or two other labs are doing similar work today; we really have a unique niche in the world."

ECT recently helped his group identify the way in which the cholera-causing bacterium *Vibrio cholerae* kills its intestinal competition, the common bacterium *Escherichia coli*, by delivering a toxin. By imaging the cholera cells, the researchers discovered tubes inside the bacteria: sometimes they were long, skinny, and filled with toxin; at other times, they were short, wide, and empty.

"Using fluorescence light microscopy and ECT, we were able to discover that these tubes are outer sheaths assembled around an inner javelin-like rod. When a cholera cell bumps into an *E. coli* cell, the outer sheath contracts and propels the inner rod of toxin through a port in the cholera cell's own membrane," explains Jensen. "The rod then punctures the *E. coli* cell and delivers the toxin. We call that a spring-loaded molecular dagger." Jensen says this discovery may lead to using this structure for medical purposes, such as designing entirely new cells to treat infection or disease.

Jensen's lab was also the first to obtain 3-D images of a complete bacterial flagellar motor. These rotary nanomachines power the miniscule whip-like flagella that are responsible for propelling bacteria through the body; their exact structure, however, had been a biological mystery until

Above: An artist's illustration shows a cholera cell injecting toxin into an E. coli *cell via a spring-loaded molecular dagger (in orange.) Grant Jensen's lab at Caltech was the first to discover this phenomenon using sophisticated imaging techniques.*

Jensen used ECT to unveil the details of how they are assembled. Thanks to that technological peephole, he and his colleagues were able to visualize how—like an outboard motor anchored to a boat—the flagellar motor in a bacterium is held in place on the outside of the cell membrane by a structure that remains fixed, allowing the flagellum itself to spin and move the cell through liquid.

One of Jensen's other significant ECT-powered advances in structural biology is the development of a better picture of how some key HIV structures form; determining the 3-D arrangement of layers present at different stages of the virus's development has helped us understand how certain antiviral drugs block HIV's growth and proliferation, he says. His lab has also helped quiet the debate over the existence of cytoskeletons in bacteria—an idea that had previously been dismissed by researchers who simply did not have powerful enough imaging technologies to see the microbes' intracellular filament skeletons.

"We've made a lot of progress in understanding why bacteria and viruses have the shapes that they do, which is key to learning about how they work," Jensen says.

Jensen compares ECT to taking apart a car; you may know nothing about cars, but if you pull an engine apart to look at all of its pieces and how they fit together, you can at least begin to intuit how it works.

"What satisfies me most is understanding how biological processes work at a mechanical level," says Jensen. "Taking pictures of the cells can, at times, advance our understanding faster than any other approach."

Visualizing Space and Time

For chemist Ahmed Zewail, however, "taking pictures" isn't quite enough. Zewail's goal is not only to visualize biological specimens—or, in his case, any kind of molecule or nanostructure—at their most basic level, but also to see the details of the fundamental

reactions they undergo … in real time.

He tackled the time element first, pioneering the field of femtochemistry, the study of chemical reactions occurring at the timescale of the femtosecond, which is one-millionth-of-a-billionth of a second. For these innovative efforts, he was awarded the 1999 Nobel Prize in Chemistry. Less than 10 years later, Zewail revolutionized the field yet again by building on his early femtochemistry work to create the four-dimensional (4D) electron microscope, which reveals objects not only in the usual three dimensions, but incorporates time into the image as well.

Most electron microscopes use a steady stream of electrons for illumination; Zewail's new technology, on the other hand, employs the precision-timed release of individual electrons—doled out one by one—to produce images of objects at the atomic scale. Each electron contributes to a picture representing a still at a given point in time. Like the frames in a film, the sequential images generated by many millions of such electrons can be assembled into a digital movie of motion at the atomic scale.

"We used to use laser light to probe and interrogate chemical systems, but that has been supplanted by the electron because the electron has properties that make it very convenient for looking at extremely small things," says Spencer Baskin, a senior scientist who has worked in Zewail's lab for 24 years.

Above: A DNA nanostructure as seen through the 4D electron microscope invented by chemist Ahmed Zewail at Caltech's Physical Biology Center for Ultrafast Science and Technology.

Resolving the details of such "extremely small things" is impossible to do with laser light, since its wavelengths are much larger than any given nanostructure. To get a "picture" of something, you need a wavelength that's at least on the same scale or, preferably, smaller than the object you are trying to resolve. Electrons are perfect for this task because their wavelengths shrink as their velocity increases; thus, they can be accelerated so that their wavelengths are a picometer, or a trillionth of a meter, in length, making it possible for researchers to capture the details of nanostructures at very high resolutions. By sending individual electrons out at known intervals of time, the 4D microscope takes that process to the next level, not only capturing tiny objects but also tracking their precise motions in real time.

"What satisfies me most is understanding how biological processes work at a mechanical level. Taking pictures of the cells can, at times, advance our understanding faster than any other approach."

"Electron microscopes have been around for a long time, but the 4D microscope extended the technology to the time domain, so we can resolve very fast processes that couldn't be seen before," says Ulrich Lorenz, a postdoctoral scholar in Zewail's lab. "This is where the power of this technique comes in: combining time and spatial resolution to get completely new information about how processes evolve."

The 4D microscope was invented at Caltech's Physical Biology Center for Ultrafast Science and Technology (UST). The center is directed by Zewail, who created it in 2005 to investigate the fundamental physics of chemical and biological phenomena. Thus far, UST researchers have been able to use the 4D microscope to study a wide range of processes in different materials, including thermal expansion and phase transitions, chemical bond dynamics, and nanomechanical motions. The latter studies yield information about mechanical properties such as stiffness, which is of particular importance for maintaining strength and integrity in everything from colossal edifices to the tiniest of nanoscale structures.

Such insights are similarly important when talking about biological structures; knowing the mechanical properties of fabrications made of DNA, for instance, is crucial to building sturdy biotech tools, such as DNA nanotubes for drug delivery in the body. And, indeed, earlier this year, Zewail and his group reported on a breakthrough 4D experiment that began with DNA stretched over a hole embedded in a thin carbon film.

They cut away several DNA filaments from the carbon film to create a three-dimensional, free-standing arrangement under the 4D microscope. Next, the scientists set the DNA strands vibrating using heat generated by a laser and imaged those motions using electron pulses. By determining the frequency of the DNA's oscillations, they were able to directly measure their construction's stiffness, a measurement that had never before been possible.

Recently, using a similar approach, the group also studied the material properties of protein assemblies called amyloids, which are believed to play a role in many neurodegenerative diseases.

"Because we're using an instrument that had never been built

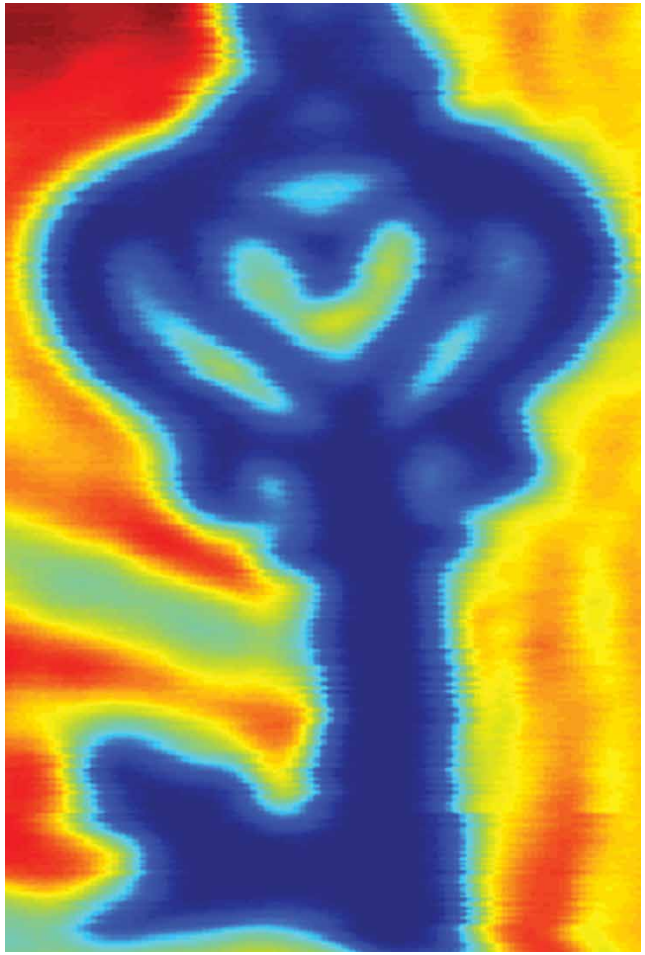

Above: A terahertz image shows a key inside an envelope; without the use of this technique, the object would be invisible to the naked eye.

before, everything we look at is something new that no one has ever clearly seen in motion before," Baskin says. "The difficulty is in finding the things that are going to be extremely important."

The end goal, he notes, is to be able to see individual molecules and atoms in the process of rearranging, taking a series of snapshots of where the atoms in a molecule are at each point in time as it goes through a reaction.

"People have frequently called our end goal a molecular movie," says Lorenz. "If you could watch chemistry—this rearrangement of atoms to make new things and transform compounds—in a movie, then you would really have it all. That would be incredible."

Transparent Technology

Engineer Ali Hajimiri isn't as interested in looking at things, per se, as in looking *through* them. It's a sideline of the work his lab has done for the past 15 years, exploring solutions for a host of different problems in electrical engineering. Five years ago, he began to work on silicon microchips able to generate and radiate high-frequency electromagnetic waves, called terahertz (THz) waves. These waves fall into a largely untapped region of the electromagnetic spectrum and were being used to penetrate numerous materials and render image details in high resolution—but only in systems that were generally too bulky and expensive for widespread use.

"There were clearly potential applications for these terahertz waves—like advanced security operations and medical diagnostic purposes," says Hajimiri. "We saw it as a big challenge worthy of our efforts, so we decided to see how far we could get with a terahertz system—how small we could make it."

In December 2012, Hajimiri's team published a paper announcing that they had developed THz imaging chips that used the high-frequency waves to see through—and image items cloaked under—materials such as fabric, plastic, paper, and wood. While the THz waves work much like X-rays, they do not carry enough energy to remove electrons from atoms or molecules and so don't create the same ionizing damage. In addition to providing images of things that are normally hidden, the system can also use spectroscopic data to detect the chemical signatures of things like chemical weapons, illegal drugs, and explosives.

"What we did was find new ways of making complete THz systems using an integrated-circuit chip-manufacturing process similar to those used to make image sensors for cell phones," says Hajimiri. "Done in large volumes, it's extremely low cost—approximately a dollar per chip set. Therefore, THz chips could be used extensively across multiple platforms ranging from cell phones to computers to other handheld devices. I think they could be ubiquitous in the long run."

In addition to being cheap, the new chip set—in which one chip acts as a light source and the other as a detector, or camera—is no bigger than a fingertip and sends out signals that are more than a thousand times stronger than those possible using existing approaches, Hajimiri says.

He and his lab members were able to use the THz chip set to detect objects hidden inside all kinds of items; for instance, they've been able to recognize a bullet stashed inside a teddy bear and reveal a razor blade in a piece of plastic.

"The first time we saw such an image—a snapshot literally representing years of effort—we were jumping up and down with excitement," says Hajimiri. "It's one thing to create a device or component that you think could be used to do certain things in the future. It's a completely different thing to actually see it working."

The beauty of the technology, Hajimiri says, is that the system can be adjusted and dynamically controlled. If, for instance, you want to image the inside of something that's soft and that, therefore, the waves can see through more easily, you can operate at a lower power. If you need to see through a much more complex or dense object, you can just crank up the power.

The technology has seemingly limitless potential applications, Hajimiri says. A THz scanner could look into large packages, crates, or even machinery. Various industries could use the scanner to check equipment for defects without having to take the object apart. For gaming or human-machine interfaces in general, THz systems could have even more impressive implications. Since the technology can be used to track movement rather than as an imaging technique, people can use it to communicate with their computers from across the room through certain gestures or even eye movements, says Hajimiri.

"Current human-machine interactive gaming systems like Kinect for Xbox are really responding to big movements of the limbs," he notes. "With terahertz waves, a gaming system would be able to detect the slightest movements of the eye, track where a user is looking, monitor their breathing, or even detect a heartbeat—it can detect even such very small displacements."

Hajimiri thinks this terahertz technology could even be used for medical applications, such as searching for tumors inside the body noninvasively, with just the wave of a handheld scanner.

"There are always new applications that we haven't thought about yet," he says. "Imagine having that sensor on your phone, and just sliding your phone across something to see what's inside it. I find all the possibilities very exciting and challenging."

"Everything we look at is something new that no one has ever clearly seen in motion before."

Indeed, Hajimiri says, the only limits in imaging technology—as in most of science and engineering—are researchers' imaginations, not their abilities.

"To use the old cliché, we need to be able to really think outside of the box and come up with new ways of doing things instead of succumbing to preconceived notions about what can or cannot be done," he says.

And when it comes to getting the first glimpse at something nobody has ever seen before, Caltech's imaging researchers are not only thinking outside the box, they're building there, too.

After all, as Lorenz points out, "Seeing something that you normally wouldn't be able to see is just so cool. It's why people do science!" E&S

John S. (Spencer) Baskin is a senior scientist in chemistry and chemical engineering.

Ali Hajimiri is the Thomas G. Myers Professor of Electrical Engineering.

Grant Jensen is professor of biology and an investigator with the Howard Hughes Medical Institute. HHMI, the National Institutes of Health, the Beckman Institute, and gifts from the Gordon and Betty Moore Foundation help support his ECT work.

Ulrich Lorenz is a postdoctoral scholar in chemistry.

Ahmed Zewail is the Linus Pauling Professor of Chemistry and professor of physics.
His lab's 4D electron microscopy studies are sponsored through grants from the Gordon and Betty Moore Foundation, the National Science Foundation, and the Air Force Office of Scientific Research.

What's the Big Deal about Big Data?

For businesses from high-tech companies to banks, big data can mean big money. For Caltech researchers, however, this new glut of information means a new way of doing science.

by Marcus Y. Woo

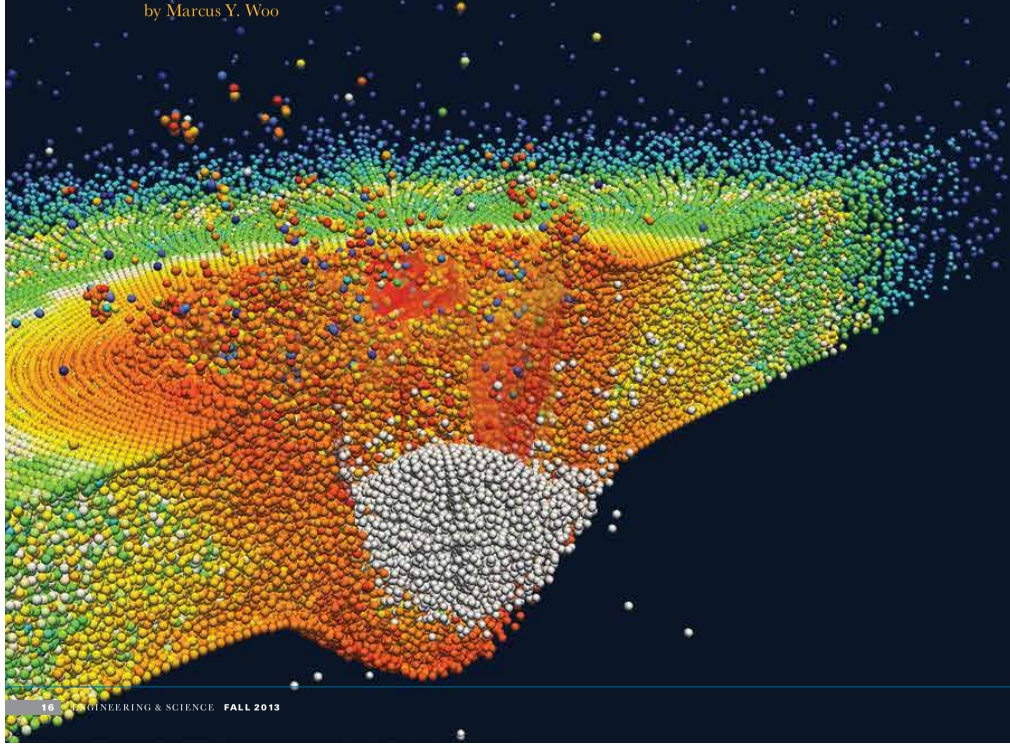

THERE'S NO DOUBT that we're awash in more information than at any other point in history. Every time you swipe a credit card, buy something online, do an Internet search, or upload a photo or video, you add to the global flood of data. Thanks to the exponential rise of cheaper and faster computers—as described by Moore's law, in which Gordon Moore (PhD '54) accurately predicted that the number of transistors on a computer chip would double every two years—we can now collect, process, and store more data than we know what to do with. Such facts and figures as stock market fluctuations, financial loan information, people's "likes" on Facebook, and their shopping habits are potential gold mines—but only if these numbers can be turned into tangible, useful knowledge.

Of course, targeted advertising and potential corporate profits are just the tip of the information iceberg. Data is inundating all aspects of society; some experts say we are on the cusp of a transformative shift. We know from experience that we can mine these unprecedented heaps of information to glean insights into everything from medicine to the environment. The Human Genome Project's analysis of all the genes in our DNA, for instance, has revealed the genetic factors that predispose certain people to particular diseases, leading to better diagnoses and treatments. Scientists' constant monitoring of the earth is helping us understand how our climate is changing and how we can best respond to other hazards, like earthquakes and mudslides. Even President Obama's reelection campaign used large amounts of data, combined with sophisticated statistical analysis, to target potential voters like never before, a tactic that's been credited with his victory.

Which may or may not be why, last year, President Obama announced a $200 million Big Data Research and Development Initiative to improve the ways we take advantage of and learn from massive data sets in such areas as health care, the environment, national defense, and education.

Despite its name, when it comes to big data, size (or, as researchers tend to refer to it, volume) isn't everything. There's also velocity and variety—which, added to volume, form the so-called three Vs. After all, the huge quantities of data are being produced, collected, and disseminated so rapidly—a few gigabytes of measurements per second from the Higgs-boson-finding Large Hadron Collider, for instance—that scientists and engineers need to continually create new computer algorithms and techniques to be able to sort the important information from the useless.

In addition, there are often so many different kinds of data involved in studying a single problem that it becomes a real challenge to integrate it all and extract any kind of coherent insight. To monitor the global climate, for instance, scientists need to keep track not only of local temperatures but of sea and ice levels and the presence or absence of a multitude of greenhouse gases to gain an understanding of the system as a whole.

It's this level of complexity that distinguishes big data from the data of the past. While data has been getting bigger for decades, it has now become so abundant, complex, and rich that its underlying meaning is not always self-evident, and conventional approaches to understanding it no longer suffice.

That bigness is changing many areas of science, such as astronomy. Instead of measuring a specific thing, be it a gene or a single galaxy, scientists now grab data on *everything*—the whole genome or large swaths of sky—and only later comb through it for potential discoveries. "There are things you can do now that you couldn't do without this data," says astronomer George Djorgovski. "Data complexity—that's the really interesting part. It's where the new, exciting things happen."

Still, big data isn't just going to deliver scientific breakthroughs on a silver platter. Big data may help answer questions, says molecular biologist Barbara Wold, "but big data itself isn't an answer. It's not magic—nor should anyone expect it to be magic." Hard work, a little ingenuity, and the scientific method will always apply, Wold says.

But the big-data craze isn't all hype, says Mark Stalzer, former executive director of Caltech's Center for Advanced Computing Research; not by a long shot. "There's an underlying truth to it," he says. "There has to be something there. Actually, I think there's a lot of great stuff there."

The next few pages describe not only the impact big data is having on the science going on at Caltech, but also how Caltech computer scientists and engineers are creating the techniques and infrastructure that will be needed for us to navigate our data-intensive future. What follows is in no way a comprehensive look at big-data science at Caltech; there's simply too much of it. Instead, as Stalzer says of the field itself, "We're just scratching the surface."

Left: A frame from a simulation of a ballistic impact, looking at the stress experienced by the more than one million particles involved when a steel spherical projectile 1.778 mm in diameter hits a 1.6 mm-thick aluminum alloy plate at a speed of 2.7 kilometers per second.

Seismic Networks

The Southern California Seismic Network (SCSN)—run by Caltech and the U.S. Geological Survey—operates seismometers at more than 400 sites spanning the region. These sensors monitor every quiver and shake in the ground, recording seismic activity and sending the data via microwave signals, satellite, and the Internet to Caltech, where it's stored and analyzed. Earthquakes are immediately and automatically identified, located, and given a magnitude.

The data rates aren't so big as to be a problem—yet. But they will be a challenge as the SCSN researchers continue to build up the network, says Caltech geophysics professor Robert Clayton. This year alone will see the addition of 100 more sensors, which will add considerably more data for the seismologists to juggle. One possible solution is to upload the data to the cloud—which would also remove the inherent problem of having a data center located in the middle of earthquake country.

But it's not only the SCSN that's a data goldmine: Caltech also operates the Community Seismic Network (CSN), a denser network of about 300 sensors designed for the home or office, centered in the Pasadena region. Data from these sensors is automatically processed and uploaded to the cloud. The eventual goal is to have at least one sensor on almost every block—as well as in schools, hospitals, and on each floor of high rises. That would require expanding the network by a factor of over 100, says Mani Chandy, Simon Ramo Professor and professor of computer science at Caltech, and it would also require new computer architectures to process the increased data flow, in which several thousand signals must be processed every second. "Current algorithms aren't fast enough to deal with that amount of data," he says.

A Greener Cloud

"The cloud" sounds like somewhere magical—an ethereal place where much of our computing happens and where our email, photos, and music live, along with a seemingly endless amount of other information. But in reality the cloud isn't quite so perfect.

The cloud consists of massive data centers—enormous warehouses containing thousands and thousands of computers, humming away 24/7. Companies like Google, Microsoft, and Amazon operate tens of thousands of these data centers all around the country. And because the computers need to be on at all times—delays or interruptions are bad for business—they carry a substantial environmental cost.

"These data centers are huge power sucks," says Caltech professor of computer science Adam Wierman, who's working to make such centers environmentally sustainable. Accounting for 2 to 3 percent of the nation's energy use, data centers emit as much carbon as the airline industry, he says. An investigative report last fall by the *New York Times* found that some data centers waste at least 90 percent of the electricity they take from the power grid. And many are in violation of various environmental regulations.

To help remedy the problem, engineers are working to develop more efficient hardware—such as processors that can run at higher temperatures and require less cooling—and some data centers are starting to run on renewable energy. But renewable energy is unpredictable: it's not always sunny, and the wind doesn't always blow. That's where Wierman comes in.

Data centers are managed by software that determines which server should do what when. To help these centers deal with erratic energy sources, Wierman and his colleagues have developed new algorithms that optimize how the centers are used. Say you want to watch a movie online. Instead of having your network access the movie through the nearest data center, even if it happens to be cloudy there, the new algorithms would send the task to a center in sunny Arizona, where solar energy is available. Or, if one data center is unusually busy, the algorithms would then distribute tasks to other data centers that happen to be underused at that time.

A large fraction of a data center's tasks involve backing up data or doing updates and other jobs that don't need to be completed right away. The new algorithms therefore delay nonurgent jobs while prioritizing those that require immediate attention. And if certain servers aren't needed at all at a particular time—say in the middle of the night, when demand is lower—then they will be shut off.

Although companies tend to get nervous when you start shuttling tasks around and turning servers on and off, the researchers have showed—on a fundamental, theoretical level—that their algorithms are indeed reliable and will save companies money in the long run. "We've been able to give really rigorous guarantees on the algorithms," Wierman says. Although the vast majority of data centers have yet to adopt these sustainable approaches, Wierman is partnering with Hewlett-Packard, which supplies server systems to other companies—including Apple—to implement the algorithms.

Wierman is now beginning to apply his algorithms to the integration of renewable-energy-powered data centers into the electrical power grid itself. For instance, when there's high demand on the grid—say on a sweltering summer day—a utility company could pay a renewable-energy-powered data center to lower its energy usage by, for instance, delaying nonurgent computational tasks. The result would be more available energy to be used elsewhere on the grid during those peak times. By providing such a power boost, the data centers would act like a battery that has stored away extra energy to inject into the grid, Wierman explains. "It's a huge win for the grid because batteries are expensive. It's going to be a long time before large-scale batteries are widely available." Using his algorithms in this way, he adds, will hopefully propel us ever closer to a sustainable future.

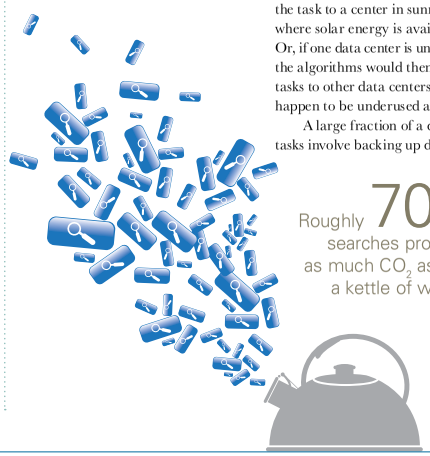

Roughly 70 Google searches produce as much CO_2 as boiling a kettle of water.

Image Search

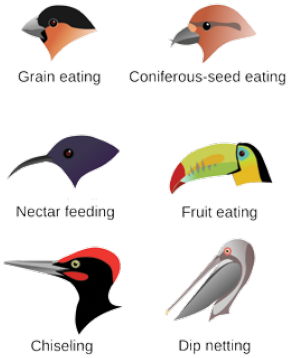

Images are among the richest forms of data—and the web is full of them. To search for a particular image, computer algorithms target key words associated with the desired picture. But what if you want to look up something whose name you don't know but that you can picture in your head—you know, that broken thingamajig in your car engine or that colorful bird in your backyard?

In most cases, images are not adequately cross-referenced, linked, or indexed to make such a search possible, explains Pietro Perona, Caltech's Allen E. Puckett Professor of Electrical Engineering. Such images are like the mysterious dark matter that pervades the universe—they're everywhere, yet invisible. "Images account for the largest portion of the web's data, and we don't know how to treat them," he says. "This is big data to the tenth power."

To address that problem, Perona and his colleagues are working with a group at UC San Diego to develop a visual encyclopedia that combines image-processing and machine-learning algorithms with expert crowdsourcing. They've dubbed it Visipedia because, in the same way that Wikipedia relies on the public for content, it relies on both experts and regular users to submit, label, and annotate images.

The researchers are starting relatively small, building Visipedia around bird images so as to take advantage of the enthusiasm and dedication of bird-watchers. The idea of Visipedia is for you to be able to upload a picture of a bird you've never seen before and get back more pictures of the same species—as well as a Wikipedia-like entry that identifies and describes it.

Initially, humans will do most of the image annotation. But as the annotators develop a systematic way to describe the images and identify important features—like the shape of a beak—the researchers hope that their computer-vision and machine-learning software will be able to learn enough from the human annotators to eventually do the same job itself. "You can think of this as a network of people and machines who are collaborating to achieve a certain goal," Perona says.

Ultimately, the Visipedia the scientists envision will be almost entirely automated—and it won't be aimed only at ornithologists and bird-watchers. This type of image recognition and searchability would be great for online shopping, Perona notes, or as a powerful tool for use in science, which is becoming increasingly dominated by visual data, organizing and automatically searching images ranging from stars and galaxies to cells and tissues and even atoms. Such a visual encyclopedia could help scientists hunt for images of a specific kind of rock or geological feature, for instance, from amid the thousands of pictures taken of the martian surface. "Once we get it just right for birds," Perona says, "you can imagine an enormous number of situations in which you would want something of this sort."

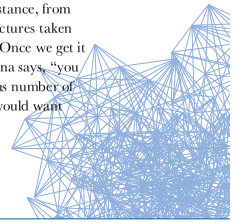

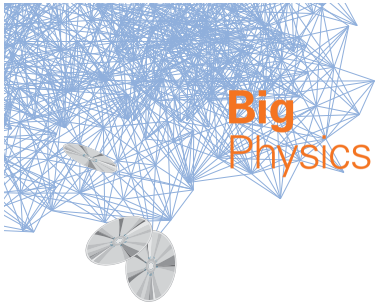

Big Physics

The Large Hadron Collider (LHC) at the European Organization for Nuclear Research (CERN) in Geneva—where last summer physicists discovered the Higgs boson—is a machine capable of slamming trillions of protons together at up to 99.9999991 percent of the speed of light, so fast that a proton careens around the accelerator's 27-kilometer circular track 11,000 times in one second.

The accelerator creates 600 million collisions per second, generating a flurry of other particles, which then decay into yet more particles. Detectors like the Compact Muon Solenoid—used by Caltech physicists at CERN—measure the velocity, position, electric charge, mass, and energy of every particle. That's a lot of data.

Indeed, there's so much data that sharing it with the thousands of physicists worldwide poses quite a challenge. In the 1990s, when construction of the LHC began, Caltech professor of physics Harvey Newman and computational scientist Julian Bunn came up with

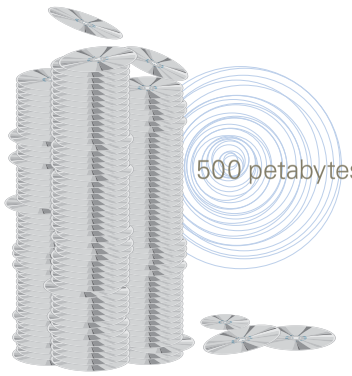

500 petabytes ≈ 100,000,000 DVDs

what turned out to be the best solution: a tiered system through which different types of data trickle down from CERN to institutions around the world for storage and sharing. With this distributed system, the data can be accessed by all the LHC partners without everything needing to be copied and sent to everyone.

CERN is the sole Tier 0 institution that creates and stores all of the raw data. There are 13 Tier 1 institutions that store the data for different regions of the world; that information is then distributed to and analyzed at hundreds of Tier 2 and Tier 3 institutions. About 500 petabytes of data are stored among all the institutions worldwide.

But that's not the whole of it. There's also computer-simulation data, the quantity of which is 10 times greater than that of experimental data, Bunn says. Comparing the simulation data with the experimental data allows physicists to search for unexpected phenomena—any discrepancies that might herald a new particle or even a new kind of physics. This simulation data, too, is distributed via the LHC grid, but because anyone can generate such simulations, it can be sent as easily from lower to higher tiers as from higher to lower.

The group led by Newman and Caltech physics professor Maria Spiropulu has also developed ever-better data-transfer methods; thanks to their work, more than 250 petabytes were transmitted through the LHC computing grid in 2012 alone. As the LHC continues to crank up its collision rates, Newman says, the flood of data will reach the exabyte range (an exabyte is a billion gigabytes). Over the next 10 years, data volumes and transfer rates are expected to grow a hundredfold.

A Faster Internet

A decade ago, it was impossible to transfer the large data sets that are now common at the Large Hadron Collider (LHC). The problem lay with the computers' so-called protocols—the systems of rules that dictate how data is transferred throughout a network like the Internet, setting up connections and automatically resending information that gets lost or delayed.

While the protocols of the early 2000s were satisfactory for the Internet and most people's needs, physicists like Caltech's Harvey Newman knew they would not be able to handle the oncoming deluge of LHC data. To solve the problem, Newman teamed up with Steven Low, professor of computer science and electrical engineering, and one of Caltech's experts on information networks.

At the time, Low says, there was no systematic way to design a protocol that would work for the huge networks required by physicists. "So what we did was try to really understand the problem by stepping back and developing a mathematical model of such networks," he explains. Working with professor John Doyle and a group of electrical engineers and computer scientists at Caltech, Low developed a deeper, structural understanding of these networks that allowed the team to build a protocol that could be as big and complex as needed. "This was not possible before," he says.

Low's ideas led to a new protocol called Fast TCP (for Transmission Control Protocol), which Newman used to set a new data-transfer record each year from 2003 to 2008. Newman and his group have since developed a sophisticated application called Fast Data Transfer—whose principal author is Caltech computational scientist Iosif Legrand—which doesn't just establish a protocol but optimizes the way huge data sets are transferred across the world, which is essential for doing LHC physics. This has allowed the team to continue breaking records: last fall, they hit a record-setting 339 gigabits per second, which is equivalent to sending one million full-length movies in one day.

Demands for high-speed data transfer have continued to grow, even among nonphysicists. And so, in 2006, Low and his colleagues started a company called FastSoft to commercialize Fast TCP. Last year, FastSoft was acquired by Akamai Technologies, which helps everyone from NBC to NASA deliver their online content. So that Grumpy Cat video you watched the other day? There's a good chance it was brought to you by Fast TCP.

Biology Gains Perspective

Barbara Wold, Bren Professor of Molecular Biology at Caltech, wants to untangle the complex webs of interacting genes—called gene regulatory networks—that determine whether a cell will ultimately become a muscle cell, a bone cell, or some other part of the adult organism. These regulatory networks consist of genes that turn one another on and off; how those interactions play out determines the cell's fate.

To understand this process in full genetic and biochemical detail, Wold and her colleagues use genomics, a field that focuses on the genome, with its 20,000 genes and hundreds of thousands of newly mapped regulatory elements. She says that the availability of "genome-wide" data for humans and key model organisms is transforming how scientists approach many problems in modern biology.

Of course, the amount of information is large, because each human or mouse genome consists of 6 billion DNA bases. This calls for new data-mining tools and ways to visualize data. Currently, integrating thousands of genomic datasets to extract new relationships among genes and their regulatory elements is a major challenge.

In June, Wold coauthored a white paper that announced a new global alliance for sharing genomic and clinical data. The ultimate goal of the alliance—of which Caltech is one of 70 founding institutional partners—is to pool data, including genome data and information about treatment received and outcomes. The resulting vast reservoir of data will help authorized researchers discover new causal relationships, doctors make better diagnoses, and scientists formulate new hypotheses. This is expected to be especially powerful for genomic diseases of high complexity such as cancer and autism.

"Creating and using large public databases that draw on data is a style of basic biology research that began with the first genome sequences," Wold says. "It is very exciting to see it fusing with clinical medicine to the benefit of both."

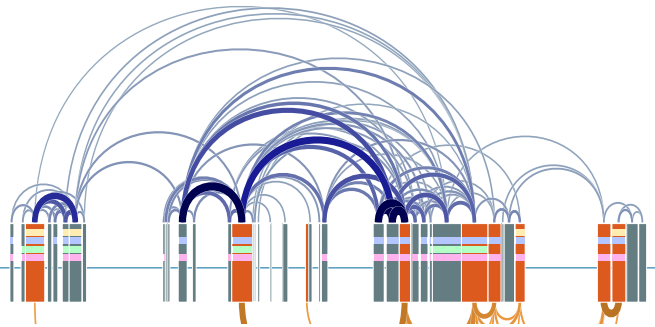

Left: Within a genome are sequences of DNA that regulate other genes, dictating when they turn on or off. Sometimes, those regulatory sequences and the genes they control are in completely different sections of the DNA strand—separated by thousands or even millions of DNA base pairs. This diagram illustrates the physical interactions between the genes in the mouse genome that are needed to turn a precursor cell called a myoblast into a myocyte, a type of skeletal muscle cell.

Better Networks

Mobile-phone base stations, fiber-optic cables, and satellites allow us to reach almost anyone anywhere on the planet. Our communication networks are huge, linking together millions of components—such as cell-phone towers and internet servers. In our data-driven world, the size and complexity of these networks will only increase.

The problem is, engineers don't have a systematic way of designing such intricate networks to function in the most efficient manner possible. "Network design is more of an art form than a science," says Michelle Effros, the George Van Osdol Professor of Electrical Engineering at Caltech and an expert on network and information theory. "People get good at it through experience and intuition."

Network components work differently linked together than they do as individual devices, researchers have discovered in recent years. But without a way to rigorously predict how a network will behave as a whole, engineers have to resort to trial and error, resulting in inefficient networks that can slow traffic and decrease reliability.

Now, Effros and her colleagues have devised some new mathematical models of generic network components that *do* predict how they would work when pieced together in a network. "Proving that such modeling is even possible is a surprising result," she says.

The researchers have so far developed models for the five most fundamental network components—which can be used to analyze all networks built from these components. Effros's ultimate goal is to keep enlarging this library, integrating the models into a piece of software that others can then use to design any kind of network they want. Using this tool, network engineers could, for instance, compare and optimize designs on a computer before they actually begin construction.

"The same ideas apply no matter what your network is," Effros says—whether you're talking about wireless networks, the Internet, or the sensors some grocery stores have installed on their shelves to monitor the freshness of foods. Some researchers are even exploring the possibility that such models can be used to understand the genetic networks that govern the development of embryos, she says.

"If we don't figure out how to use our networks properly and design them better, the path we're on will be limited," she says. "To keep expanding our communication capabilities requires real advances."

A **Numbers** Game

Analytics—also known as advanced statistics—has changed the way pros play baseball and basketball. And now it's changing Caltech basketball as well.

What is advanced statistics? "It's information beyond the box score," explains Oliver Eslinger, the head coach of the Caltech men's basketball team. Knowing how many points someone scored in a game isn't as informative as *how* that player got those points: when and where on the court he made his shots; whether those shots were contested jumpers or easy layups; whether they were the result of set plays or were assisted and, if so, which teammate passed the ball.

Eslinger and his coaching staff record and annotate every detail of every game and practice, developing formulas to better quantify the performance of each player and the team as a whole. Every play provides a plethora of data that can be used by coaches during practices, when determining game-day strategies and lineups, and when instructing players on how they can improve during the off-season.

The use of advanced statistics is still rare in Division III, Eslinger says. "I'd be surprised if any other DIII coaches are doing what we're doing," he notes. "I'd like us to be at the forefront of analytics in college sports." And for a place like Caltech, that's certainly fitting.

Flashes in the Night

Most stars remain static over the course of a human lifetime, but some rapidly brighten, dim, flare, or even explode. Objects whose brightness varies significantly are called transients, and they help astronomers understand how stars live and die, and how black holes form.

Caltech astronomers using Caltech's Palomar Transient Factory (PTF) and the Catalina Real-Time Transient Survey (CRTS) have trained automated telescopes on the heavens to search for such objects—and have brought in quite a haul. They have discovered thousands of variable stars and bright, black-hole-powered galactic nuclei. But the surveys may have made their biggest mark by recording thousands of new stellar explosions called supernovae.

Together, PTF and CRTS have discovered almost 4,000 of the more than 6,300 supernovae known so far. And the surveys' numbers rise daily. By collecting so much data, astronomers have been able to discover entirely new kinds of astronomical objects—such as new classes of supernovae—prompting new scientific inquiries, says astronomy professor George Djorgovski. And PTF and CRTS are only the beginning.

The Zwicky Transient Facility, also led by Caltech, is a PTF upgrade that's set to begin its survey in 2015. And, within the next decade, the Large Synoptic Survey Telescope (LSST) will begin a constant watch of the night sky with greater sensitivity and resolution than ever before.

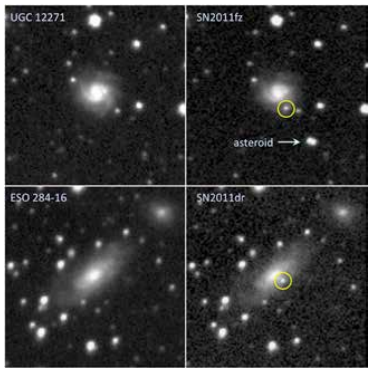

Above: Examples of discoveries from the CRTS survey show, on the left, galaxies before supernova explosions. The images on the right show the galaxies with the supernovae circled. When the top right photograph was taken, there was an asteroid passing through the field at the same time; the software identifies asteroids and separates them from astrophysical transients like supernovae.

While PTF and CRTS might detect a few tens of transients per night, Djorgovski says, LSST should be able to find as many as 10 *million*.

In addition, this fall Caltech's Owens Valley Radio Observatory Long-Wavelength Array is set to begin imaging the entire viewable sky every second to search for transient signals at radio frequencies—in particular, signals from nearby exoplanets. These signals arise when particles spewing from the planets' stars interact with the planets' magnetic fields. During its hunt for transient signals, the array will generate 2.5 gigabytes of raw data every second, a rate similar in scale to that of the Large Hadron Collider, says astronomer Gregg Hallinan, who's leading the radio transient search.

The impending explosion of data makes the development of tools capable of analyzing all of it increasingly important, Djorgovski says. That's why he and his colleagues at Caltech and JPL are developing basic pattern-recognition and machine-learning algorithms to identify the information that's worth keeping.

Still, no matter how good the tools are, there is simply too much data for professional astronomers to analyze. The solution, Djorgovski says, is to make that information accessible to all. "When you have little data, data is precious," he says. "But when there's so much data that you can't possibly do it yourself, it's actually irresponsible *not* to let others do it."

The CRTS already makes its entire data set publicly available; the hope is that amateur enthusiasts will dive in and make their own discoveries. According to Djorgovski, it's this democratization of science that will be the most important consequence of the data explosion in astronomy.

"Anybody with an Internet connection has the same opportunity as astronomers at Caltech," he says. "And that's great." E&S

NAVIGATING THE BRAIN'S MYSTERIES

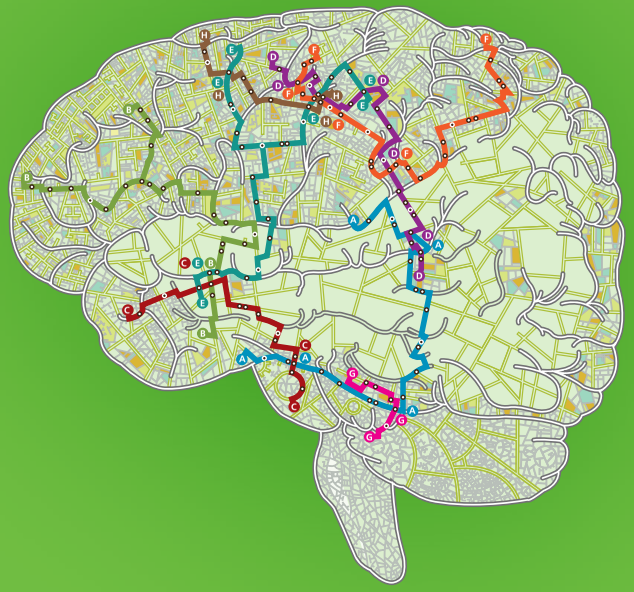

CALTECH RESEARCHERS (AND THE WHITE HOUSE) say it's time to piece together a dynamic map of the brain—one that shows its complex trafficking across trillions of neuronal connections. Addressing this grand challenge could just be the technological moon shot of a generation.

By Kimm Fesenmaier

One morning this past April, nanoscientist Michael Roukes and neurobiologist Thanos Siapas sat amid the scientists and engineers packed into the East Room of the White House. They listened as Francis Collins, director of the National Institutes of Health (NIH), introduced President Barack Obama as "our scientist in chief," and watched as the president took his place behind the podium.

After acknowledging the attendees as "some of the smartest people in the country"—and joking about the questionable appropriateness of his new scientific title—the president got down to the business of the morning: outlining "the next great American project," the BRAIN (Brain Research through Advancing Innovative Neurotechnologies) Initiative.

"As humans, we can identify galaxies light-years away, we can study particles smaller than an atom," the president said. "But we still haven't unlocked the mystery of the three pounds of matter that sits between our ears."

Obama went on to describe how his newly proposed initiative would aim to change that. This newest of his Grand Challenges would be to obtain not only a thorough map of the brain and its roughly 100 billion neurons, but a dynamic picture of how that complex organ works in real time. Ultimately, scientists could use this knowledge to pick apart how we think, learn, and remember as well as how to better treat disorders such as schizophrenia, Parkinson's disease, post-traumatic stress disorder, and Alzheimer's disease.

Roukes, Siapas, and many of their colleagues had long awaited an announcement like this one—an acknowledgment of the project they had championed for several years, that of a large-scale effort that would focus top scientists and engineers on revealing the details of the brain.

LET'S BACKTRACK

It all really started taking off for Roukes when, in September 2010, he traveled to Oslo as the director of the Kavli Nanoscience Institute at Caltech to attend the annual Kavli Prize symposium. He struck up a conversation there with the directors of several Kavli neuroscience centers about the maturation of technologies coming out of nanoscience and their potential to improve our understanding of the brain.

Why was Roukes, a physicist and nanoscientist, even thinking about neuroscience? That dates back to a decision made more than 15 years before by former Caltech provost Steve Koonin. In the late 1990s, Gilles Laurent (now a director of the Max Planck Institute for Brain Research) told Koonin that one of the things his field was sorely lacking was the ability to insert tiny electrodes into the brain that would allow scientists to record signals from multiple neurons at the same time. That got Koonin thinking.

"I knew that Michael was expert at fabricating very tiny things," Koonin says. "So I connected Gilles and Michael up and provided a bit of seed money to grease the interaction. My intuition in doing so was nothing more than having two accomplished faculty interested in reaching out across disciplinary boundaries, and knowing that new instrumentation almost always leads to new science. Who knew it would blossom into what it did? Such are the rare pleasures of academic administration."

What blossomed was a close friendship between Roukes and Laurent, and a lasting collaboration at the intersection of neuroscience and nanoscience. After a couple of small pilot projects, Roukes helped Laurent introduce the use of tiny neural probes that could be mass-produced.

In 2002, Siapas joined the Caltech faculty as a neurobiologist interested in brain circuits and the functions of memory and learning. These complex functions arise as a result of the coordinated activation of large populations of neurons distributed throughout the brain. In order to elucidate these brain patterns, Siapas wanted to capture large-scale recordings from freely behaving animals. His interests aligned perfectly with those of Roukes and Laurent, and the three began thinking about ways to enhance the scale and quality of electrophysiological recordings as well as develop prototype devices to explore different research directions.

As a result of these interactions, Roukes says he became increasingly fascinated by the brain. "In fact, I've sort of switched the center of my activity toward biological applications of nanotechnology," he says, "with the brain now being a principal effort. And all of this is the result of Steve Koonin saying, 'Get to know this guy, Gilles. See what happens.'"

BAM

A year after Roukes's informal conversation in Oslo with the neuroscience directors, a symposium took place outside of London, hosted by the Kavli Foundation, the Allen Institute for Brain Science, and the Gatsby Charitable Foundation. There, a

number of participants from neuroscience and nanotechnology—including Roukes, Siapas, and Caltech neuroscientist David Anderson—came together to identify new opportunities for technological development at the interface between these disciplines. "There was spirited discussion about the best avenues and approaches to take," Anderson says.

During the course of the symposium, a subgroup of participants—including Roukes; George Church, one of the leaders of the Human Genome Project; Rafael Yuste, a neurobiologist from Columbia University; and others—got together and began to formulate a new project.

Capitalizing on the momentum from that symposium, Roukes wrote a "technical foundations document" along with Church and Paul Alivisatos, a nanomaterials scientist and director of the Lawrence Berkeley National Laboratory. The document laid out a roadmap of sorts, describing the nanotechnologies that would need to be developed to fuel a neuroscience revolution. In June 2012, a group including Roukes, Church, Alivisatos, Yuste, and two others published a paper in the journal *Neuron* describing what they dubbed the Brain Activity Map (BAM) Project. They followed that up earlier this year with a brief overview of the project, this time in the journal *Science*.

One of the central points in all of these documents was that many brain functions may emerge as a result of neuronal activities taking place in physically separate regions of the brain at the same time. Monitoring such disparate activities is no easy task; neuroscientists today are generally restricted to using electrodes that allow them to study brain activity from one neuron at a time up to only a few. But brain circuits involve millions of such nerve cells, each with thousands of connections that might be rearranging all the time. Focusing on individual neurons could lead researchers to miss the forest for the trees.

On the other hand, imaging technologies such as functional MRI and magnetoencephalography (MEG) are able to capture whole-brain activity—but at the expense of single-cell specificity. They allow researchers to see which brain regions are activated while a subject participates in a particular activity, for example, but provide little detail in terms of which neurons are involved, how they're connected, and under which circumstances they fire. In other words, that forest is looking pretty nice, but what happened to the trees?

To get at the elusive middle ground, where researchers would be able to image, understand, and eventually manipulate collections of neurons at the level of brain circuits, the BAM advocates called for the development of new tools that would allow them to record every activity spike from every neuron in a circuit. Current imaging techniques cannot record activity from enough neurons and do not reach sufficient depths within the brain tissue to achieve this goal. Nor are current techniques for gathering electrophysiological measurements able to record activity from enough neurons in dense enough patches. However, the authors argued, there are promising research avenues that could improve the situation in each of these areas. They also suggested that entirely new methods for wirelessly, noninvasively recording neuronal activity could prove useful. For example, they wrote in the *Neuron* paper that they think "it will ultimately become feasible to deploy small wireless microcircuits, untethered in living brains, for direct monitoring of neuronal activity."

To take steps toward that goal, Roukes and Siapas have started a Beckman Institute pilot project at Caltech, in which they are developing arrays of tiny electrodes called nanoprobes that would be able to measure brain activity from far greater numbers of neurons than is currently possible. Typically, to record electrical signals from neurons, researchers insert hand-assembled bundles of four small wires into the brain tissue. Siapas and others have managed to get recordings from as many as 25 of these bundles at once, but it has proved difficult to scale up beyond that.

Roukes and Siapas's research would allow them to mass-manufacture tiny silicon probes that could record neuronal activity from denser populations of neurons by using many recording sites along the lengths of each probe. "Using these techniques, we can make a new generation of needles that are much finer and have many, many more recording sites," Roukes says.

THE DETAILS

If nothing else, the published articles and less formal white papers produced by the proponents of the Brain Activity Map Project sparked a conversation in Washington, D.C., about the potential benefits that could be realized by a large-scale, brain-related national initiative. In February, President Obama even alluded to such a project in his State of the Union address—two months before announcing the BRAIN Initiative. "Every dollar we invested to map the human genome returned $140 to our economy," he said. "Today our scientists are mapping the human brain to unlock the answers to Alzheimer's … Now is not the time to gut these job-creating investments in science and innovation. Now is the time to reach a level of research and development not seen since the height of the Space Race."

The president would later propose getting the BRAIN Initiative started with a budget of about $100 million for fiscal year 2014—with funds coming from the National Institutes of Health, DARPA (the Defense Advanced Research Projects Agency), and the National Science Foundation. Several private organizations, including the Kavli Foundation, have also said they will contribute to the effort. But details about the allocation of funding, the areas of research, and the initiative's goals and milestones are still to be defined.

To that end, the NIH has appointed a high-level working group of neuroscientists to review available information, to recommend goals that are in line with the vision of the initiative, and to come up with a scientific plan for achieving those goals.

Caltech's Anderson is a member of that 15-person group, which NSF director Collins refers to as the BRAIN Initiative's "dream team." The team's charge? First, compile a list of research areas tagged for immediate funding. Then submit a full report in June 2014.

Anderson considers himself privileged to be part of the working group. And even more to the point, he says, "I'm thrilled that the president of the United States has recognized the importance of understanding brain function."

Asked about the group's progress, Anderson notes that he and his BRAIN colleagues are just beginning deliberations. "Our plan," he says, "is to solicit input from a broad range of scientists. I think it will be fascinating and instructive to listen to the different voices in the neuroscience community and see what kind of consensus can be reached."

While many have applauded the initiative and its ambitious scope, there are detractors who worry, among other things, that the project's funding will steal from other neuroscience projects or that the brain is the wrong subject for such a focused project. Anderson, however, sees the BRAIN Initiative as "an exciting opportunity to accelerate progress in our understanding of brain function in health and disease by promoting new technology development and applications."

Roukes agrees. "This is an incredible opportunity to do a moon shot in terms of the technology that will be developed, which will democratize how the next generation of neuroscience is done," he says. "We would be foolish not to capitalize on this moment." ES

David Anderson is the Seymour Benzer Professor of Biology and an investigator with the Howard Hughes Medical Institute.

Yaser Abu-Mostafa is a professor of electrical engineering and computer science. His work is supported by Intellectual Ventures.

Michael Roukes is the Robert M. Abbey Professor of Physics, Applied Physics, and Bioengineering. His neuro/nano work is funded by a National Institutes of Health (NIH) Director's Pioneer Award and grants from the G. Harold & Leila Y. Mathers Foundation and the National Science Foundation (NSF).

Thanos Siapas is a professor of computation and neural systems. He receives funding from the Gordon and Betty Moore Foundation, the Mathers Foundation, the NSF, and the NIH.

Learning From Machines

Neuroscience and nanoscience are not the only fields that are likely to benefit from the new **BRAIN Initiative**. Certainly, trying to understand how hundreds of thousands or millions of neurons connect and behave across a range of timescales will produce what has been called a "data deluge"—and managing and making sense of that flood of information will require new tools for data management and analysis.

A Kavli Futures Symposium was held at Caltech in January, bringing together 16 scientists from a range of fields to discuss this issue. The organizers estimated that recording from a million neurons a thousand times per second would generate 100 terabytes (102,400 gigabytes) of data per day. However, if you compressed the data, you might be able to get down to about 3,000 terabytes of data per *year*—a huge amount of information, but not beyond the range of other big-data projects.

Machine-learning expert Yaser Abu-Mostafa participated in the symposium. "The magnitude of the project and the amount of data will be completely impossible to handle," he says, "unless you have a method that will be able to do the needed mapping, interpretation, and analysis in an automated way." Abu-Mostafa believes that such a method will come from machine learning—an area in which, he adds, Caltech has a lot to contribute.

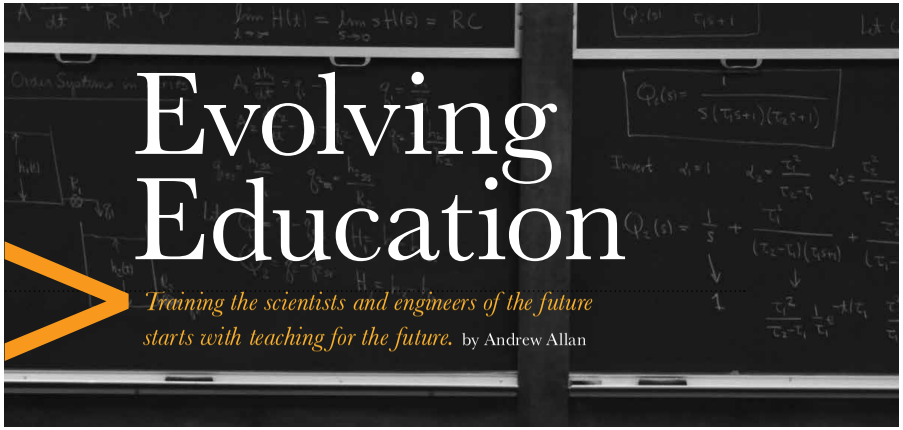

Evolving Education

Training the scientists and engineers of the future starts with teaching for the future. by Andrew Allan

As any evolutionary biologist will tell you, if you don't evolve, you perish. Throughout time, this has proved true not just for entire species, but for societies, organizations, companies, nations, and—especially—institutions of higher education, where change is essential to reaching the widest variety of students and to remaining competitive in an ever-expanding field of educational options.

And so, by modernizing the undergraduate core curriculum, restructuring the campus's writing center, founding a dedicated center for teaching and learning, and using the online arena to expand the Institute's educational reach, Caltech is undergoing an educational evolution—an educational rewiring, if you will—that is changing the way students connect with faculty, faculty connect with students, and both interact with the information and ideas they're encountering together.

"One size fits all doesn't work anymore. We have a diversity of academic programs and a diversity of students, and the old way just isn't the solution," says Vice Provost Melany Hunt.

"There's a renewed vigor in the way we think about undergraduate education, and a commitment among the faculty to improve it," adds Jonathan Katz, chair of the Division of the Humanities and Social Sciences, who has been closely involved in Caltech's recent initiative to institute a series of changes to update its core curriculum. This sequence of general education requirements has not been revised for nearly two decades. "These changes have important implications not only for how our students will learn," Katz says, "but for how we as faculty will teach them."

In the midst of all this reworking and rewiring, Minnesota photographer and educator Martin Springborg came to Caltech as part of his work on a nationwide photographic essay he's compiling; he spent three days taking shots of faculty and students as they went about the business of teaching and learning. The photographs on the next few pages are from that look at a day in Caltech's educational life.

Cassandra Volpe Horii (left), director of the Caltech Center for Teaching, Learning, and Outreach, works with Professor of Biology and Geobiology Dianne Newman (on left at head of table, below) and her Principles of Biology course TAs to provide guidance in implementing the latest teaching methods, developing better lectures and more targeted homework assignments and exams, and obtaining and analyzing feedback from students.

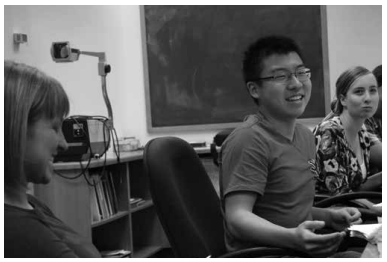

At the Hixon Writing Center—under the leadership of campus writing coordinator and lecturer in writing Susanne Hall—students work one-on-one with professional and peer tutors to generate ideas, develop arguments, organize their thoughts, and enhance clarity in academic, technical, and personal writing.

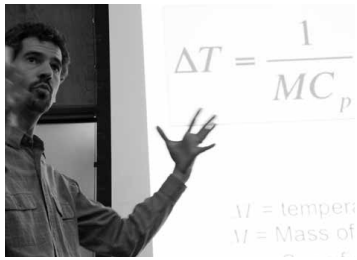

Left: Caltech has a long history of exceptional teaching by faculty such as Professor of Geology and Geochemistry Paul Asimow, winner of the 2012 Richard P. Feynman Prize for Excellence in Teaching, who here is seen teaching freshmen in his geology survey course, Earth and Environment.

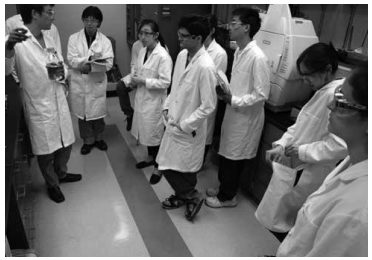

Right: Teaching and learning at Caltech take place as much outside the classroom as inside. Students taught by David Tirrell, Ross McCollum-William H. Corcoran Professor and professor of chemistry and chemical engineering, discuss their latest research in his Biomolecular Engineering Laboratory, which focuses on the design, construction, and characterization of engineered biological systems.

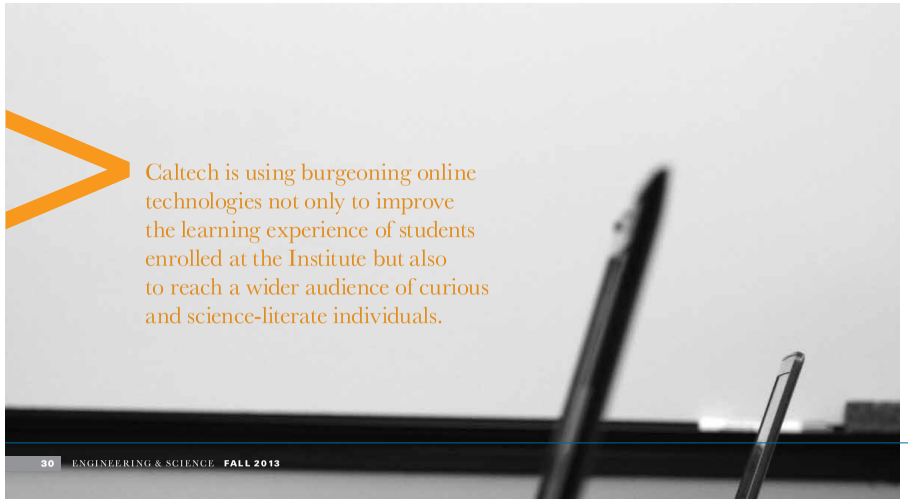

Caltech is using burgeoning online technologies not only to improve the learning experience of students enrolled at the Institute but also to reach a wider audience of curious and science-literate individuals.

The online arena is changing the way professors impart knowledge to students—both in university settings around the world and at home. Caltech is using burgeoning online technologies not only to improve the learning experience of students enrolled at the Institute but also to reach a wider audience of curious and science-literate individuals.

In mid-2012, Caltech joined with Coursera, an online education company that offers courses from the world's top universities and organizations free of charge. Through Coursera, Caltech launched a series of MOOCs—massive open online courses—at the end of last year. Bren Professor of Biology Henry Lester, Professor of Astronomy George Djorgovski, and Professor of Economics and Neuroscience Antonio Rangel were the first Caltech professors to teach MOOCs through the Coursera platform.

"MOOCs provide the latest technology for several of our goals, such as developing Caltech students' communication skills and pedagogical experience," Lester says. "These activities will help to develop scientific professional careers in the 21st century."

Vice Provost Hunt says that MOOCs have virtually unlimited potential for impacting teaching and learning, not only allowing Caltech faculty to share their research with a global audience but also enhancing education on campus by allowing students to go back and review lectures or brush up on previous lessons. Adding an online component also allows teachers to rethink the structure of their classroom presentations.

"We hope that more and more of the passive process of listening to lectures happens at home on the student's laptop," says Hunt, "so that the classroom experience can become a more engaging, interactive process."

It is with the goal of improving teaching and learning at Caltech—and ultimately the Caltech experience—that these types of additional resources are being made available to both students and teachers. This educational evolution, which is changing the way students and faculty connect and engage with each other and with information, will ensure that Caltech remains a uniquely challenging environment. E&S

Game Theory, Austen Style

Could Jane Austen have been a social scientist? Michael Chwe (BS '85) thinks so.

In his book *Jane Austen, Game Theorist*, Caltech alumnus Michael Chwe—now an professor of political science at UCLA—makes the case that the beloved author was a careful observer of strategic thinking.

Chwe first became intrigued by the idea while watching the film *Clueless*, the 1995 adaptation of Austen's novel *Emma*.

"I never took a literature class at Caltech, much to my detriment," Chwe laughs. "But when I read *Emma*, I was struck by how Austen carefully crafted scenarios that mirrored modern strategic thinking."

Chwe contends that more than 100 years before the mathematician John von Neumann established game theory as a discipline for studying conflict and cooperation, Austen was already on the case.

Take *Emma*'s Emma Woodhouse, who tries to play matchmaker between the simple Harriet Smith and Mr. Elton, the village vicar. Emma's attempts at manipulation backfire when Mr. Elton turns out to have intentions toward Emma herself. In this case, Chwe says, Emma falls into a classic trap—overconfidence in her own strategic powers.

And that, Chwe believes, is the point.

"Austen's novels do not simply provide 'case material' for the game theorist to analyze," Chwe writes in his book, calling the novels as a whole "an ambitious theoretical project, with insights not yet superseded by modern social science."

Austen famously questioned the barriers between social classes. But with Chwe's help, she has now crossed another divide—the one between art and science.

Michael Chwe (center) gets a chance to meet two of Jane Austen's characters thanks to TACIT (Theater Arts at Caltech) and two Caltech graduate students. Emma Woodhouse, right, is played by **Meg Rosenburg**, who is earning her PhD in planetary science, and Mr. Elton is **Kelvin Bates**, who is working toward a PhD in environmental chemistry.

Sorting and Sharing a Mountain of Gene Sequences

John Quackenbush (BS '83)

Since the first announcement in 2000 that a working draft of the sequence of the human genome had been assembled, gene-sequencing technology has sprinted forward. It's now possible to have your chromosomes read in less than a day and for the price of a used car.

But once researchers have a genome in hand—what then? That's what fascinates John Quackenbush.

"The challenge is no longer how to generate these vast amounts of sequencing data," says Quackenbush, a professor of computational biology and bioinformatics at Harvard. "Now we need to effectively collect, manage, and interpret the information in ways that make it useful and—above all—private and secure."

In 2011, Quackenbush founded GenoSpace to create an advanced cloud-computing environment that enables complex genomic and clinical data to be securely stored, analyzed, and made accessible to a broad range of researchers and health-care providers.

In a ceremony at the White House this past summer, Quackenbush and 12 other scientists were recognized as Open Science Champions of Change for their commitment to sharing scientific data.

"With these new collaborative tools," Quackenbush says, "we can accelerate the pace of scientific progress while respecting the important role each individual patient can play in the process of medical discovery."

Big Data for the Masses

Rajiv Ghanta (BS '05) and
Cesar Del Solar (BS '04, MS '05)

While looking for a way to track his personal health and financial data on a single screen, Rajiv Ghanta realized that he had an information problem—there was too much of it.

"Caltech prepared me to be comfortable around numbers," says Ghanta. "But the largest system in the brain is the visual cortex. Our brains are optimized to handle larger amounts of information when it is presented visually, making it easier to detect patterns."

So in 2010, Ghanta and fellow Caltech graduate Cesar Del Solar founded Leftronic, a web service that allows customers to aggregate data in real time from a wide variety of sources and then present that data in an elegant dashboard with enough charts, maps, and other statistical eye candy to turn the casual number tracker into a true data wonk. The dashboards can then be displayed on a website, an HDTV in the lobby of a building, a smartphone, or any web-enabled device.

"We started Leftronic because we felt everyone should have the tools to understand their data quickly and easily," Del Solar says.

Now you can have a whole new stat to keep track of: the number of hours you spend staring at your nifty new Leftronic dashboard.

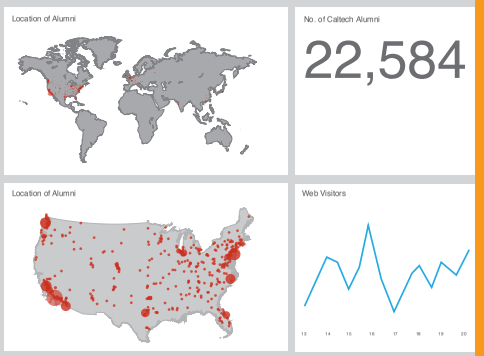

Alumni stories provided by the Caltech Alumni Association. For more about these stories and to read about other alumni in the news, go to *alumni.caltech.edu*.

Francis H. Clauser, 1913–2013

Francis H. Clauser (BS '34, MS '35, PhD '37), the Clark Blanchard Millikan Professor of Engineering, Emeritus, passed away on March 3 at age 99. Born in the decade following the Wright Brothers' first powered flight, he was a founder of modern aeronautics and helped usher in the Space Age.

Upon earning bachelor's degrees in physics at Caltech, Clauser and his twin brother, Milton, entered the aeronautics program run by Theodore von Kármán. After receiving their doctorates, the brothers joined the Douglas Aircraft Company in Santa Monica, where Francis soon became the director of aerodynamic design research, assembling a team that profoundly influenced aviation design by developing mathematical methods for shaping tails, wings, engine nacelles, and air scoops.

In 1946, Clauser founded the aeronautics department at Johns Hopkins University. He chaired the department until 1960, recruiting leaders in the field from many countries to create a world-class research center.

In 1965, Clauser was invited to the newly created University of California campus in Santa Cruz to set up the engineering school, and in 1969 he returned to Caltech to chair the Division of Engineering and Applied Science. He stepped down in 1974 but remained the Millikan Professor of Engineering until his retirement in 1980.

Clauser was a Fellow of the American Institute of Aeronautics and Astronautics, the American Physical Society, and the American Association for the Advancement of Science. He was also a member of the National Academy of Engineering, the scientific research society Sigma Xi, the engineering honor society Tau Beta Pi, and Caltech's Gnome Club. He became a Distinguished Alumnus in 1966, one of the initial 23 people to be honored.

Clauser is survived by his sister, Betty Celeste Valois; his son, Wolf laureate in physics John Francis Clauser (BS '64); and his daughter, Caroline Helen Ryan.

To learn more about Francis Clauser's life, visit caltech.edu/content/francis-clauser.

Donald Coles, 1924–2013

Donald Coles (MS '48, PhD '53), professor of aeronautics, emeritus, passed away on May 2. He was 89 years old.

After receiving his undergraduate degree from the University of Minnesota, Coles came to Caltech for his graduate degrees; he then spent his entire research career at the Institute, retiring in 1996. An expert on the properties of turbulent flow and the dynamics of rotating fluids, Coles provided the first reliable experimental data for local surface friction in supersonic flow. Among his teaching achievements was the development of a course in technical fluid mechanics that bridged the gap between classical shear flows and their applications in the design of control valves and jet flaps.

Coles was a fellow of the American Institute of Aeronautics and Astronautics, the American Physical Society, and the American Association for the Advancement of Science, and he was elected to the National Academy of Engineering in 1984. In 2000, the Donald Coles Prize in Aeronautics was established at Caltech; it is awarded at commencement to the aeronautics PhD "whose thesis displays the best design of an experiment or the best design for a piece of experimental equipment." In 2011, GALCIT (the Graduate Aerospace Laboratories at the California Institute of Technology) created the Donald Coles Lectureship in Aerospace in his honor.

Coles is survived by his wife, Ellen; by his four children, Christopher, Elizabeth, Kenneth, and Janet; by his sister, Marjorie Schlaegel; and by two grandchildren.

To learn more about the Donald Coles' life, visit caltech.edu/content/donald-coles-0.

CALCULATING CALTECH: We asked alumni to provide us with an equation to sum up their Caltech experience, reminding them—of course—to define their variables. Here is what some of them came up with.

$\text{Caltech} = e^{x}\text{cellence}$

$E = mc^2$ Excellence through mathematical calculations at Caltech
Excellence equals **minds creatively challenged**

$e^{i\pi} = -1$
sums up Caltech: *rational, irrational, real, complex, whole, natural, positive, negative.*

$CIT = U/\infty$
When you go to Caltech, you create infinite possibilities for yourself.

At Caltech, $1 = 2$. Caltech requires double the concentration and effort; as a consequence, you get double the benefits.

When I graduated from Caltech, I applied for this license plate, which is the polar coordinates equation for the **number 1, which Caltech is**.

$\Delta S < 0$
where S is entropy. Although, as we know, entropy tends to increase in any closed system, locally we may assert that Caltech brings order out of disorder, knowledge out of ignorance.

My equation for Caltech is the simple exponential formula, $y = e^x$. For me, it implies the explosive amount of knowledge that continues to accumulate and accelerate.

$Caltech = e^{x}$
as in exponential accumulation of knowledge at an exponential accelerating rate.

$1 \rightarrow \infty$ As in, from the **one Caltech** comes a seemingly **infinite amount of potential and creativity**.