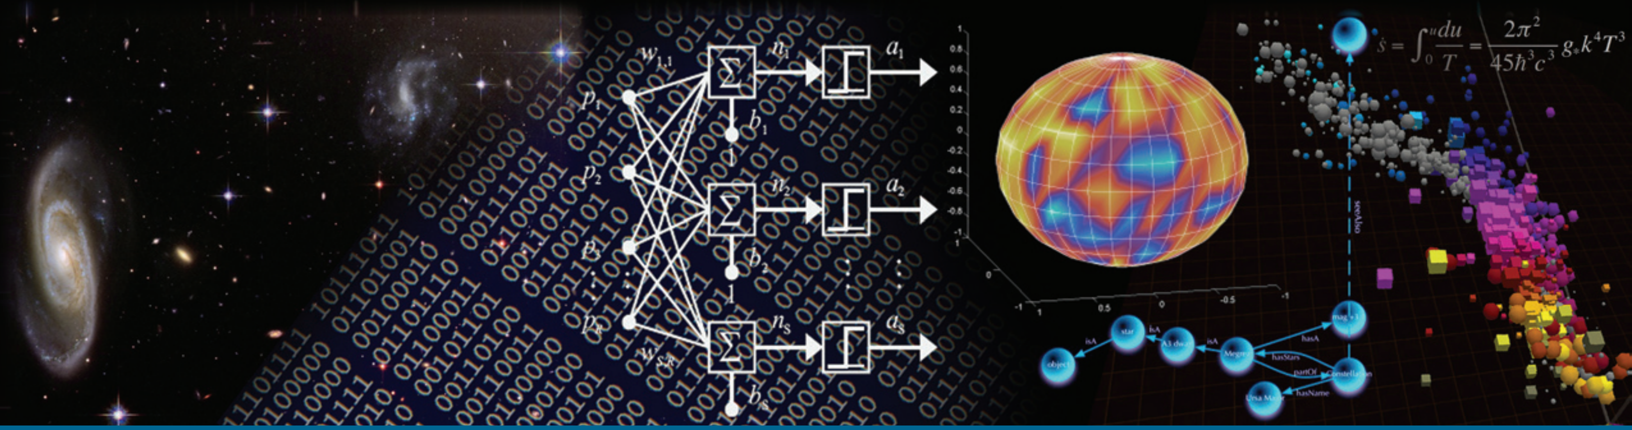

A UNIVERSE OF ASTRONOMICAL DATA

Long gone are the days of the lone astronomer perched atop a mountain, peering through an eyepiece at a smudge in the night sky. Modern astronomers, though, still have to go out and get their own data, either spending nights in the telescope's control room or observing remotely.

But for the astronomer of 2020, data will automatically come to her before she wakes up in the morning. While she sleeps, telescopes on autopilot will comb the night sky, feeding many terabytes (a thousand billion bytes) of data into computers. The computers will then mine the data, flagging any intriguing cosmic events, such as new supernovae or explosive gamma-ray bursts, for our astronomer to peruse while sipping her morning brew. Computers will also have collected information from other databases that may be relevant—multiwavelength images of that part of the sky, for example.

> "We'll be able to ask questions that we couldn't dream of asking before," says George Djorgovski.

Later, she might log onto some three-dimensional virtual environment, where virtual representations of herself and astronomers from around the globe will literally immerse themselves in the data—a colorful three-dimensional cluster of spheres that represent the hundreds of characteristics that describe, for instance, a supernova—analyzing the supernova in ways that would have been inconceivable a mere 10 years earlier.

Welcome to astroinformatics—an amalgam of astronomy and computer science. Increasingly powerful instruments, more sensitive telescopes, and ever-faster computers are inundating the astronomical community with far more data than it ever imagined. The next generation of sky surveys is expected to uncover billions of stars and galaxies and lots of other interesting things, such as supernovae, quasars, exoplanets, and possibly new objects that have yet to be discovered—each with hundreds of parameters.

The Large Synoptic Survey Telescope project, of which Caltech is a partner institution, will produce petabytes of data—roughly the amount of information contained in a billion books. Another project planned for the future, the Square Kilometre Array, is expected to yield *exa*bytes of data—a billion billion bytes. By comparison, Caltech's been involved with numerous surveys—such as the Digitized Palomar Sky Survey, Palomar-Quest, the Palomar Transient Factory, and the Two-Micron All Sky Survey—that have generated tens of terabytes of data, a factor of at least a thousand less than what future surveys will produce. And the size of data will follow Moore's Law, doubling every couple of years.

It's not just that the databases will be huge. They'll be complex—imagine trying to visualize, much less understand, hundreds parameters all at once. Sifting through the numbers is humanly impossible, so Caltech astronomers are taking the lead in developing tools to process, analyze, and understand this deluge of information.

For example, the so-called Virtual Observatory (VO), which Caltech has been a leader in developing since its inception 10 years ago, is a way to integrate all the data that's being collected by telescopes from around the world and in orbit. Every telescope's data sets are different, not only in terms of what information is gathered—Chandra is a space telescope that looks at X-rays, and Keck is an optical scope on top of Mauna Kea, for example—but also in format and how they're accessed. But with the VO, an astronomer can enter a query on the computer and get to all the relevant databases at once, without having to learn the quirks and technicalities that accompany each data set. "Federating is the word that's normally used for this," says Matthew Graham, a computational scientist at Caltech's Center for Advanced Computing Research (CACR). "You're federating these different data sets and then using online services to do things with them."

After a decade of developing the tools and infrastructure needed to get these databases to talk to each other, the project, now called the Virtual Astronomical Observatory and funded by NASA and the NSF, opened for business in May. "We're moving onto the operational phase," says Graham, a member of the program council of the VAO. "The hope is that we can

Left: The exponential rise in information is changing how astronomy is done.

Right: This image from Second Life shows two avatars immersed in a virtual world of three-dimensional data.

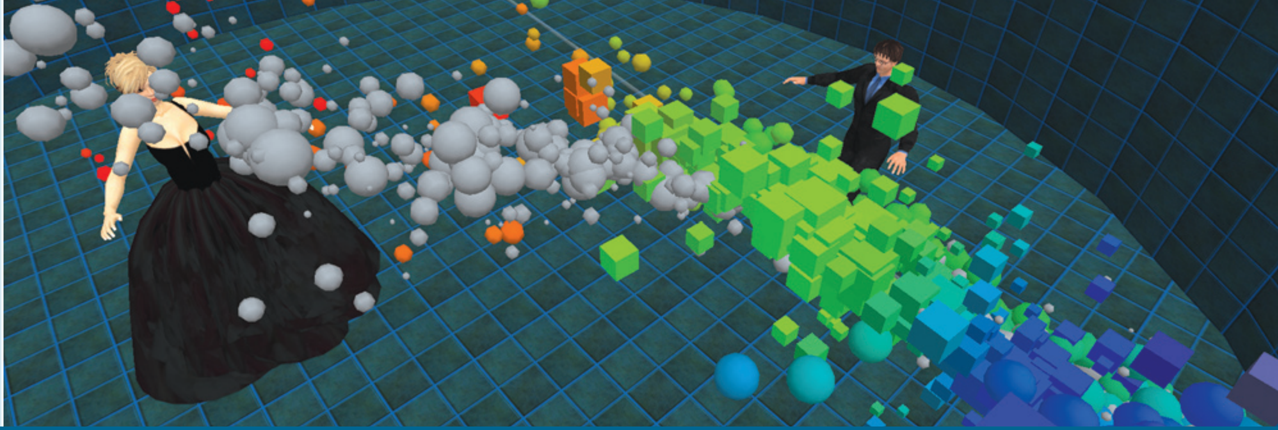

really make an impact on the community." In addition to Graham, CACR computational scientist Roy Williams (PhD '83) also plays a leading role with the VAO. Others at Caltech who are involved with astroinformatics include CACR executive director Mark Stalzer, CACR computational scientist Andrew Drake, executive director of the Infrared Processing and Analysis Center (IPAC) George Helou, IPAC scientists Joe Mazzarella and Bruce Berriman, postdoc Ciro Donalek, staff scientist Ashish Mahabal, and others.

But there's more to astroinformatics than just bigger telescopes, better computers, and more sophisticated software. "It's not just the same old stuff with more data, but genuinely new things," says George Djorgovski, professor of astronomy and principal investigator for Caltech's part of the VAO consortium. "We'll be able to ask questions that we couldn't dream of asking before, just because we didn't have the tools or the data." This past June, Djorgovski was one of the organizers of the first astroinformatics conference, held at Caltech's Cahill Center for Astronomy and Astrophysics and attracting about a hundred astronomers from around the world. Participants spent four days discussing a wide array of topics, ranging from data mining and computation to education and outreach. The conference was even broadcast live on the Web, and participants posted comments on Twitter during the talks.

The outreach goes far beyond education, as astronomers are actually soliciting the public's help. You're probably familiar with SETI@home, a program that uses your computer's downtime to search radio-telescope data for signs of extraterrestrial life. Astronomers want to take advantage of your brain, as well as your laptop. Galaxy Zoo, for example, is a website that asks users to classify thousands of individual galaxies from the Sloan Survey. Distinguishing a spiral galaxy from an elliptical one is a complex problem for a computer, but a simple one for a human. With more than 250,000 users, Galaxy Zoo has spawned similar projects to help astronomers sift through data taken by other missions, like the Lunar Reconnaissance Orbiter and the Hubble Space Telescope.

But even turning the entire world into an astronomy sweatshop will fall short. "There aren't enough humans on the planet to handle the data right now," Graham says. So researchers like him want to take this idea of "citizen astronomy" farther and figure out how citizen scientists interpret data to develop smarter data-mining algorithms. For instance, a human can identify a bright light in the spiral arm of a galaxy as a supernova, because we know that an exploding star has to live in a galaxy. Understanding this kind of contextual information is hard for a computer, but if machine-learning researchers analyze enough images in which supernovae have been spotted by humans, then some other characteristics that a computer can process may be uncovered.

This sort of technique, part of a subfield called semantic astronomy, is still in its early stages, Graham says. But a lot of astroinformatics will involve similar tools that turn the computer from a number-crunching machine into an intelligent assistant. Instead of having to mine through different databases and pick out the relevant numbers by hand, an astronomer could just type in a query in plain English—for example, "find all the data on stars within 100 light-years of us"—and the computer would cull all the relevant information from every database available, leaving the astronomer free to focus on the science.

Of course, astronomy is far from being the only field overwhelmed with information. The burgeoning field of *bio*informatics has been transforming biology for the past decade. Other sciences are facing similar challenges: real-time sensors monitoring everything from earthquakes to climate are generating a barrage of data, and Moore's Law is driving an exponential growth in information. "Any science that's using semiconductors to do detection is suddenly becoming data-intensive," says CACR's Mark Stalzer.

"Science in the 21st century is going to be different," Djorgovski adds. "The focus is shifting from having better hardware to having better software and methodology. It's going from atoms to bits to knowledge." With the smart phones, social networking, and news feeds that inundate our everyday lives, we're all experiencing a torrent of information. And like the rest of us, astronomers are learning to deal with it. —*MW*

e&s

If you want to learn more about astroinformatics, you can find slides and videos of all the talks from the Astroinformatics 2010 conference at www.astroinformatics2010.org.

TOLERATING A FETUS

Your immune system recognizes cells in your body that aren't yours, hunts them down, and kills them. So how does a fetus survive? Half of its genes and all of its tissues are unlike mom's, yet the body does not attack this invader.

Caltech biologists have discovered that a particular type of immune cell—produced in response to specific fetal antigens, proteins that stimulate the immune system—allows "pregnancy tolerance," as it's called.

"Our finding that specific T regulatory cells protect the mother is a step to learning how the mother avoids rejection of her fetus. This central biological mechanism is important for the health of both the fetus and the mother," says David Baltimore, the Millikan Professor of Biology, and recipient of the 1975 Nobel Prize in Physiology or Medicine.

Scientists had long been "hinting around at the idea that the mother's immune system makes tolerance possible," says Daniel Kahn, a visiting associate in biology at Caltech, and an assistant professor of maternal–fetal medicine at the UCLA. What they didn't have were the details of this tolerance—or proof that it was immune-related.

Now they do. Baltimore and Kahn selectively destroyed the T regulatory cells in a strain of mice bred so that all the males—including male fetuses—carry on their cells' surfaces a protein known as a "minor transplantation antigen." Female mice lack this antigen.

Normally, pregnancy tolerance would kick in and protect the male fetuses from any maternal repercussions.

So if the T regulatory cells provided the shield, their destruction would give the immune system free rein to go after the antigen-laden males—and only the males.

And indeed, fewer male fetuses survived to birth. Those that did were of significantly lower birthweight, presumably because of the inflammation caused by the mother's immune response to that single antigen.

The scientists found that pregnancy tolerance "develops actively as a consequence of pregnancy," says Kahn. "The mice are not born with it." Indeed, virgin mice showed no signs of these pregnancy-specific T regulatory cells. Conversely, the cells were found in larger numbers in mice that had given birth to male babies, with the level of T regulatory cells increasing with the number of male births.

The next step, Kahn adds, is to look at T regulatory cells and their role in pregnancy tolerance in humans—a line of research that may lead to insights into such pregnancy-related conditions as preeclampsia, in which high blood pressure and other symptoms develop in the second half of pregnancy. Preeclampsia is a major cause of maternal mortality around the world.

"There's a lot to be learned," he says. "Pregnancy is often ignored in research because it's usually successful, and because—from an immunologic standpoint—it has such complexity. Until now, it's been difficult to get a handle on how the immunology of pregnancy really works."

The work is described in an article by Baltimore and Kahn in the May 18 issue of the *Proceedings of the National Academy of Sciences*. The research was supported in part by a grant from the Skirball Foundation. —*LO* e&s

SHOOT THE MOON, HIT A ROVER

A Caltech alum has found a lunar rover that's been lost for 40 years. On November 17, 1970, Lunokhod 1 rolled off a ramp from the Russian spacecraft Luna 17 and became the first remote-controlled robot to land on another world. The eight-wheeled vehicle, about the size of a riding mower, explored Mare Imbrium, the Sea of Rains, covering 10.5 kilometers and traveling for 322 days before its handlers lost contact with it. The second Lunokhod (Russian for "moonwalker") landed on January 15, 1973, and covered 37 kilometers over four months before it overheated.

The twin Lunokhods had French-built laser reflectors on their backs, and since Lunokhod 2's exact location was known, scientists have been shooting laser pulses at it to measure the distance to the moon with extreme accuracy. On the other hand, Lunokhod 1's coordinates were known only to within five kilometers—to hit it with laser pulses, you would need to know where it is within better than 100 meters. "It's good enough to put a push-pin into a map, but not

Right: Looking kind of like a bathtub on wheels, the Lunokhod 1 is 2.3 meters long and 1.5 meters tall. The reflectors sit in the tray jutting out on the left.

Below: The final location of the Lunokhod 1 rover.

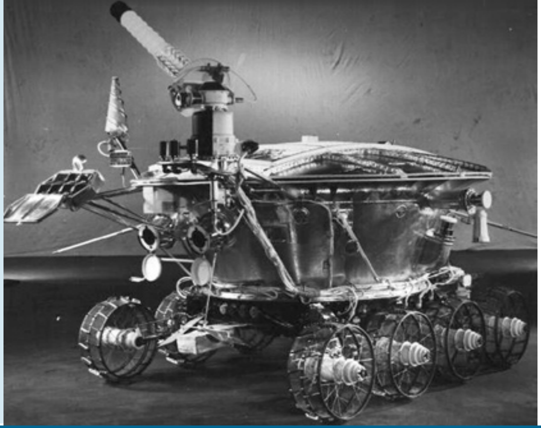

nearly good enough to make a laser search likely to succeed," says Tom Murphy (MS '97, PhD '00), associate professor of physics at UC San Diego and the principal investigator for the Apache Point Observatory Lunar Laser-ranging Operation (APOLLO).

The project is recording every tilt, tip, and wobble of the moon in order to gauge its precise orbit and test Einstein's theory of gravity, general relativity. APOLLO uses the 3.5-meter telescope at the Apache Point Observatory in New Mexico to shoot lasers at the Lunokhod reflectors or one of the other three reflectors planted on the surface by Apollos 11, 14, and 15. By timing how long it takes the laser pulse to return to the telescope, researchers know exactly how far it is to that point on the moon's surface. Each reflector consists of an array of three mutually perpendicular mirror segments arranged like the inner corner of a cube. These so-called corner reflectors bounce the light beam directly back toward its source, regardless of its direction. Apollos 11 and 14 have 100 reflectors in their arrays, Apollo 15 has 300, and the Lunokhods have 14 larger ones.

Although they had four targets, Murphy and his colleagues wanted to find Lunokhod 1 because it sits nearer to the edge of the lunar disk. The moon's axial tilt and precession cause it to wobble, and the outer part of the disk—called the limb—moves the most toward or away from Earth, which is what laser ranging measures well. Figuring out the limb's motion would thus give a more accurate measurement of the total wobble. Since the team had a rough guess as to where Lunokhod 1 was, give or take five kilometers, they hoped to hit the rover's mirrors with laser pulses. The laser beam, which leaves the telescope about three meters wide, stretches to two kilometers by the time it hits the moon. Still, nothing ever came back. "It almost seemed like a waste of telescope time," Murphy says.

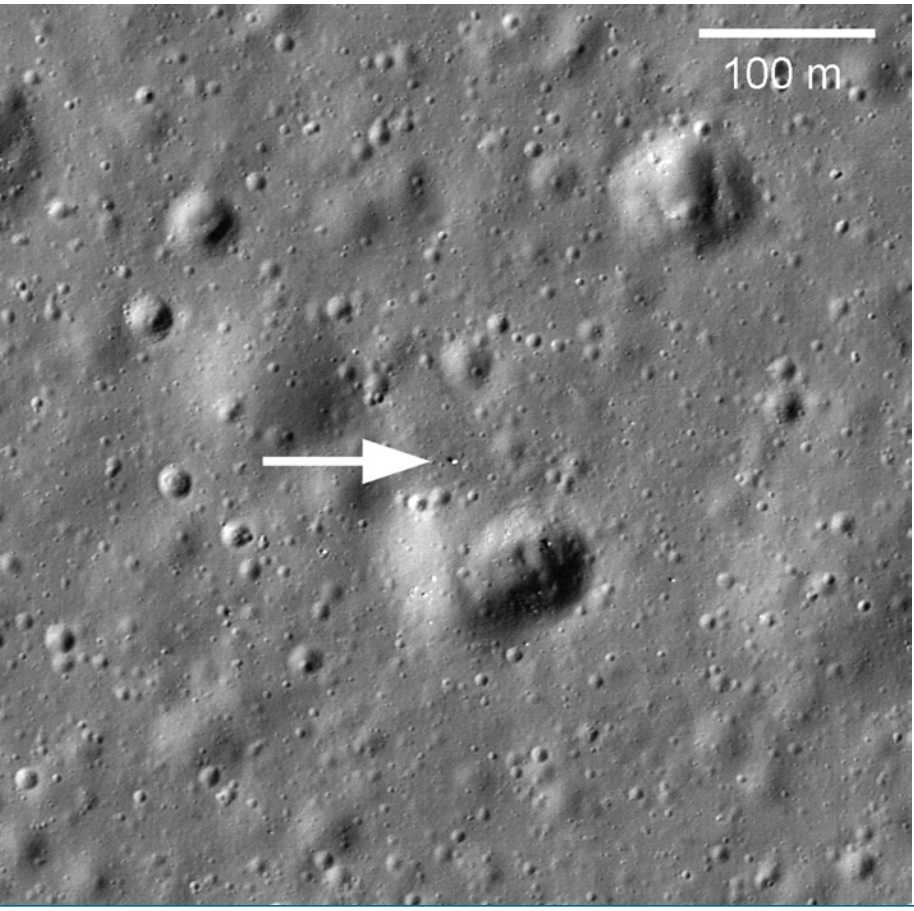

Meanwhile, NASA's Lunar Reconnaissance Orbiter (LRO) had been snapping shots, at a resolution of one meter, of all the lunar-landing sites, from those of Apollo to those of the Russian Luna missions. In March, LRO took a picture of the area around Luna 17's landing site, and spotted the missing rover. It was four kilometers from where Murphy and his team thought it was, and without LRO's help, they would've never had a chance to find it, he says. LRO was able to pinpoint the rover's coordinates to within 100 meters, giving APOLLO a target to work with.

Hitting a mirror with a reflective area of 489 square centimeters from 384,400 kilometers away is like trying to hit a grain of rice in New York City from Los Angeles. The researchers shoot 20 pulses per second, and with every pulse, they blast 10^{17} photons at the target. On a good night, maybe one photon per pulse bounces back and reaches the telescope. Because the signal is so weak, the detector has to amplify every photon it

FIRST STEPS FOR CURIOSITY

The Mars rover named Curiosity is taking shape. Engineers have installed its wheels and its remote sensing mast, which holds up the rover's set of cameras, forming its neck and head. On July 23, Curiosity took its first steps, slowly rolling across the floor at JPL. Watch a video of Curiosity's first test drive.

receives. As a result, it is only turned on for 100 nanoseconds per pulse, to avoid picking up background photons that would flood out the signal. Since light travels at almost 300,000,000 meters per second, this means the distance to Lunokhod 1 had to be known to within about 15 meters, although the team was later able to improve their technique and widen the window to 90 meters. Using LRO's altimeter, called LOLA, the researchers determined the elevation of the Sea of Rains to within five meters, and on April 22, when they started firing, photons came piling in—the first time any signal has come from Lunokhod 1 in four decades.

Lunokhod 1 was so reflective that it shocked Murphy and his team. After sending about 10,000 pulses, the team had gotten about 2,000 photons back—a bounty compared with the best-ever 750 photons that Lunokhod 2 had returned from 5,000 pulses. Overall, Lunokhod 1 is five times brighter than its twin. Since the Lunokhods are identical, and one would expect both mirrors to have endured similar degradation from dust and tiny meteorites, why the twin rovers are so different is a mystery.

The newfound rover also outshines the mirrors left by Apollo 11 and 14, making it the second brightest reflector on the moon. "Its position near the limb, combined with the fact that it's so strong, means that it will become, after Apollo 15, our most valuable target," Murphy says. APOLLO can measure distances to within a millimeter, and the team is now trying to pin down Lunokhod 1's location to that degree of accuracy. This will take about a year, after which they'll be able to get equally accurate numbers for how the moon wobbles, leading to a better fix on the center of the moon's mass and therefore the shape of its orbit. Isaac Newton's 300-year-old equations describe orbits quite well, but when you zoom in to a scale of around five meters, general relativity gives different numbers. APOLLO's data will test Einstein's theory to an accuracy of one part in 10,000.

Other than testing general relativity, knowing the precise motions of the moon will also help scientists glean details about the moon's interior structure and composition. "Imagine that you walk along a sidewalk and run into a trash can," Murphy explains. "If the can falls over, it's empty. If it wobbles, it's full of stuff." Scientists still don't know for sure whether the lunar core is entirely liquid or solid, or a combination of both, nor do they fully understand the interaction between the core and the mantle. Knowing more about lunar interior structure will tell us more about how the moon and solar system formed—and ultimately, how we all came to be. *—MW*

LUTETIA AND SATURN

Saturn hovers behind the asteroid Lutetia. The Rosetta spacecraft took this snapshot from 36,000 kilometers away with its OSIRIS narrow-angle camera. Operated by the European Space Agency, Rosetta will arrive at Comet 67P/Churyumov-Gerasimenko in 2014 and send a lander to explore the comet's surface. Researchers at JPL help with the American contribution to the mission, which includes three instruments: an ultraviolet imaging spectrometer called ALICE, a microwave instrument called MIRO, and an ion and electron sensor (IES).

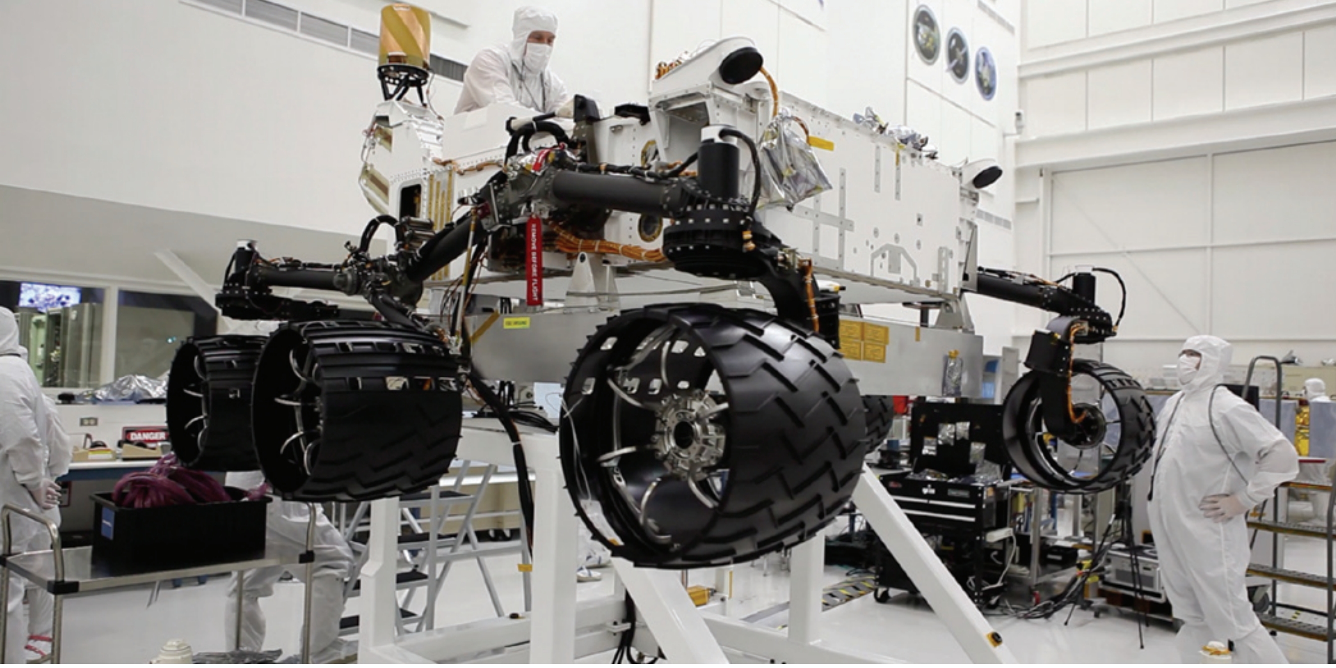

T CELLS MAKE A COMMITMENT

When does a cell choose its particular identity? That's one of the big questions in biology. We now know the answer, at least for a branch of the immune system called T cells.

The activation of a gene called *Bcl11b* is a "clean, nearly perfect indicator of when cells have decided to go on the T-cell pathway," says Ellen Rothenberg, the Ruddock Professor of Biology at Caltech.

The *Bcl11b* gene acts to shut off other genes in the stem cells from which T cells are born, allowing the stem cells to pick one of the many developmental paths open to them.

"Stem cells and their multipotent descendents follow one set of growth rules, and T cells another," says Rothenberg, "so if T-cell precursors don't give up certain stem-cell functions, bad things happen."

The conversion from T-cell precursors to actual T cells takes place in the thymus, a specialized organ located near the heart. "When the future T cells move into the thymus," Rothenberg explains, "they are expressing a variety of genes that give them the option to become other cells," such as mast cells (which are involved in allergic reactions), killer cells (which kill cells infected by viruses), and antigen-presenting cells (which help T cells recognize targeted foreign cells).

As the T cells enter the thymus, the organ sends molecular signals to the cells, directing them down the T-cell pathway. At this point, the *Bcl11b* gene gets turned on, blocking other pathways. This is critical.

"For cells that never divide again, maintaining identity is trivial. What they are at that moment is what they are forever," Rothenberg says. But T cells keep dividing as they migrate around the body and interact with other types of cells.

The Bcl11b protein "is like a switch that allows the cells to shut off stem-cell genes and other regulatory genes," Rothenberg says. "It keeps them clean—and may be necessary to 'guard' the T cell from becoming some other type of cell."

Although it is thought that many genes are involved in the process of creating and maintaining T cells, "*Bcl11b* is the only regulatory gene in the whole genome to be turned on at this stage," she adds, "and it is probably always active in all T cells. It is the most T-cell specific of all of the regulatory factors discovered so far." Among blood cells, this gene is only expressed in T cells, she says. "The gene is used in other cells in completely different types of tissue, such as brain and skin and mammary tissue, but that's how the body works. There's no confusion, because something like brain tissue and mammary tissue will never be a T cell."

When *Bcl11b* is not present—as in mice genetically altered to lack the gene—T cells "don't turn out right," Rothenberg says. Indeed, T cells in some individuals with T-cell leukemia have been found to have lost the gene. "It may make them more susceptible to the effects of radiation, because the cells don't know when to stop growing," she says. "We think that the loss of one of the two copies of the gene is enough to prevent cells from growing appropriately."

The discovery is described in "An Early T Cell Lineage Commitment Checkpoint Dependent on the Transcription Factor *Bcl11b*," a paper in the July 2 issue of *Science*—one of three papers on the *Bcl11b* gene. The paper was coauthored by Rothenberg, Caltech postdoc Long Li, and Mark Leid of Oregon State University. The work was supported by the California Institute for Regenerative Medicine, the National Institutes of Health, the Caltech–City of Hope Biomedical Research Initiative, the Louis A. Garfinkle Memorial Laboratory Fund, and the Al Sherman Foundation. *—KS* e&s

Associate Professor of Aeronautics and Bioengineering John Dabiri, shown here with his jellyfish tanks, first became interested in studying the organism's propulsion system while doing a SURF project at Caltech. Now he has his own SURF students.

SURFIN' SAFARI

Thirty-one years ago, Caltech junior Ken Libbrecht was one of 17 students in the Institute's new Summer Undergraduate Research Fellowships program, or SURF, as it quickly came to be known. At the time a unique program, SURF offered undergraduates the opportunity to pursue original, hands-on research in close collaboration with faculty mentors. The students could choose the area in which to work—a junior committed to chemistry might spend 10 weeks studying earthquakes, or a sophomore undecided about whether she *really* wanted a post-college career in the laboratory might have a better idea about it after 10 weeks of communing with a collection of petri dishes.

Whatever the neophyte researchers elected to do, the whole idea was to give them a sense of how research actually works, from that first crucial step of submitting a proposal to the final formidable one of writing a research paper. Each SURFer received a summer stipend, came to know their faculty advisers as colleagues and friends, and went on to present their research at a SURF symposium modeled on professional conferences. By the time they left Caltech, several had also put their names to one or more published articles while they were still undergraduates. Above all, they had gained a firsthand appreciation of the research experience.

Libbrecht, a physics student, spent his SURF summer working with fellow physicist and Caltech professor Steve Koonin on an aspect of nuclear theory, which resulted in a paper in the highly regarded *Physical Review Letters*. Now himself a Caltech physics professor for more than a quarter century, Libbrecht is still involved with SURF. But these days, *he's* the mentor.

Although the basic elements of SURF remain the same, quite a bit has changed since 1979, when 17 students worked with 16 mentors. This year 431 students are working with 261 mentors. The SURFers include 53 undergrads from other schools who have come to Caltech to do research; 46 who are working with scientists at Pasadena's world-renowned Jet Propulsion Laboratory; and a number of Techers who are SURFing offsite at other university campuses, national laboratories, or high-tech R&D companies.

After all this time, SURF boasts many alumni, including current Caltech professors. Two of them, John Dabiri and John Johnson, were both undergraduates at other universities when they took part in SURF, and both say the experience had a significant impact on their careers.

Back in 2000 Dabiri was an undergrad at Princeton, "and Caltech wasn't on my radar at all," he says. "I told one of my professors that I was interested in doing summer research in experimental fluid mechanics, and he suggested the names of a few professors around the U.S., including Mory Gharib, a Caltech professor in aeronautics. I had never been to California (or on a plane!), so this seemed like a good excuse."

Dabiri enjoyed that SURF summer so much that he came back to Caltech for a PhD, with Gharib as his thesis advisor. Now an associate professor of aeronautics and bioengineering, Dabiri says, "My SURF involved measurements of jellyfish swimming. I wasn't thrilled when I first heard about the project because I didn't think biology could be rigorous. But I fell in love with biological fluid mechanics, and I have been doing it ever since."

Thinking he might want to attend Caltech as a graduate student, Johnson, then an undergrad at the University of Missouri-Rolla, saw SURF as "an opportunity to learn more about life at Caltech, build up research experience, and hopefully get a letter of recommendation from a Caltech prof."

Johnson, who did his SURF with Caltech's Laser Interferometry Gravitational-Wave Observatory (LIGO) research team, has since moved on to observational astronomy, identifying and studying planets beyond our solar system. He credits his SURF experience with helping him realize that he'd rather work in a smaller research group than in a large consortium. "It taught me that I love research, but that I needed a research question of my own."

Dabiri adds, "I think SURF can be eye-opening for students who are used to classroom learning, where someone else has already solved all of the questions. In research they get to experience the frustration and exhilaration of learning something no one else knows. That certainly was my experience."

When Libbrecht returned to Caltech after receiving his PhD at Princeton, he had no doubt that he wanted to mentor SURF students himself, a sentiment echoed by Dabiri and Johnson. Still, Libbrecht acknowledges that the experience can be a bit bittersweet. Johnson was his SURF student back in '99, and, says Libbrecht, "A person feels old when your SURF students have their own SURF students!"

· · ·

This article originally appeared online in Caltech Today the week of July 13. A few days later, we received an email from Jim Morgan, Goldberger Professor of Environmental Engineering Science, Emeritus. He was curious as to whether his former student Jim Jensen had been one of those original 17 SURFers in 1979. Indeed he had. The two men have kept in touch over the years, and it turns out that both remember that first SURF summer.

"Great guy all around," says Jim Morgan of Jim Jensen. "I still remember Jim playing in the pep band. . . ." As to the importance of SURF, Morgan opines, "I believe that SURF was instrumental in stimulating his interest in a future career."

Jensen agrees, saying, "Before my SURF experience, I had no idea what research entailed. I was hooked instantly by the open-ended nature of research, by the collaboration with Jim and his graduate students, and by the small victories and seemingly enormous challenges."

Jensen now says that one of his biggest pleasures—as academic director of the University at Buffalo's Research Exploration Academy and professor in the department of civil, structural, and environmental engineering—is having "the privilege of introducing undergraduates to the joy of research."

"As I work with underclassmen in our research seminars, I often think back to those sunny summer days in Keck Lab and ask, 'What would Jim Morgan do to inspire them?'"

Jensen's 1979 SURF project looked at how metals behaved in the presence of analogs of naturally occurring organic matter. After graduating from Caltech in 1980, he went on to receive his PhD at the University of North Carolina at Chapel Hill, continuing to work on naturally occurring organics. First as a graduate student and then as a young professor, he was elated to meet some of the people who wrote the papers that had inspired his SURF project. "SURF taught me about the community of scholars I was about to join."

After that memorable summer of '79 "elbow-deep in glassware and chemicals," Jensen enjoyed another first—becoming one of the first three SURF students to make a presentation to Caltech's Board of Trustees. "The trustees were gracious and pretended not to notice my shaking knees," he says. In hindsight, he says, that SURF summer, complete with Board presentation, gave him the confidence and enthusiasm to think seriously about pursuing his own research and teaching career.

It's been a career marked by a long line of Jensen's own graduate students, numerous awards and scientific papers, and two books. And now he's working on his third: the fourth edition of *Aquatic Chemistry* by Werner Stumm and—yes—James J. Morgan. The book remains the definitive resource on the essential concepts of natural water chemistry—in fact, it's considered by many to be the field's bible.

"I was deeply humbled when Jim approached me about revising the book that he cowrote. I never dreamed that I'd be writing the fourth edition of a book that we used in his class all those years ago," says Jensen.

Adds Morgan, "And only thirty years gone by." —*PD* E&S

Former SURFer Jim Jensen is now on the faculty at the University at Buffalo.