e&S
Engineering & Science

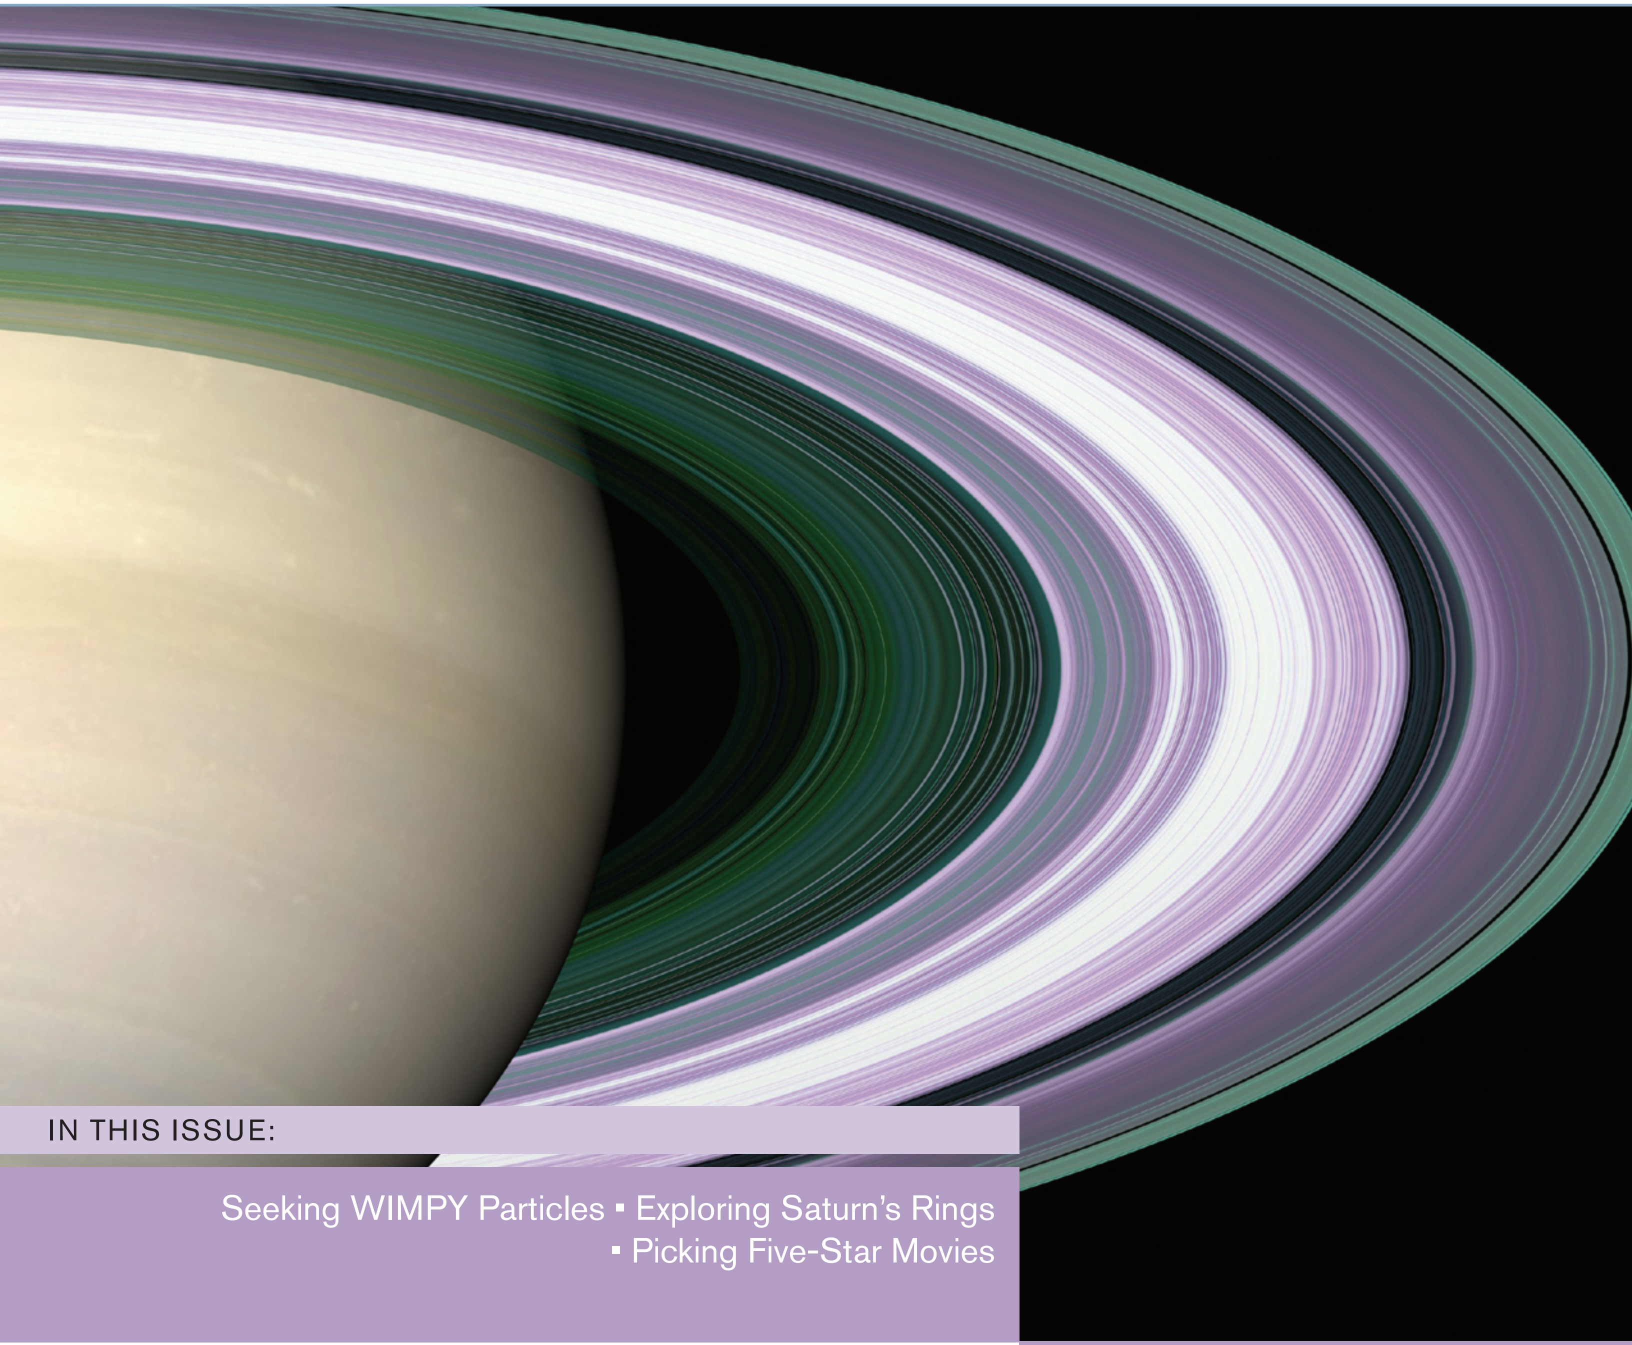

IN THIS ISSUE:

Seeking WIMPY Particles ▪ Exploring Saturn's Rings ▪ Picking Five-Star Movies

VOLUME LXXIII, NUMBER 2, SPRING 2010

California Institute of Technology

ON THE COVER

Saturn's rings are made of particles ranging from tiny pieces of dust to huge boulders. On May 3, 2005, the Cassini spacecraft sent radio signals back to Earth through the rings to measure the particles' size distribution. In this simulated image, which has a resolution of 10 kilometers, purple represents regions where all the particles are bigger than five centimeters wide. Green areas contain particles smaller than five centimeters, and blue ones have particles smaller than one centimeter. The white zones are so dense that the signal didn't get through. For more on what Cassini is showing us about the rings, go to page 22.

40 YEARS AGO IN *E&S*

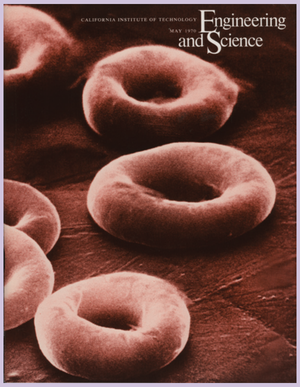

May 1970's cover shot of red blood cells was taken with a scanning electron microscope by Richard Baker, a professor at USC and a Caltech research associate. He, Professor of Engineering Science J. Harold Wayland (MS '35, PhD '37), and JPL's John Devaney were studying the cells' mechanical structure.

"The Future Isn't What It Used to Be," Arthur C. Clarke told Caltech on April 12, 1970. The dean of science-fiction writers had been invited to campus to speculate on what the world would be like in—of course!—2001. He cautioned that his conjectures were far from predictions, and indeed many of his ideas seem outlandish today. Yet the themes he touched on are still very much relevant. He recognized the inefficiency of raising animals for food, and suggested that we domesticate animals such as antelopes, tapirs, or hippos that could live off "marginal land that's of no use for anything else." We could also herd whales instead of cattle, an idea he wrote about in his 1957 novel, *The Deep Range*. And we could engineer microorganisms to transform "inedible materials—such as sawdust and wastes of various kinds—into food which we or our animals can eat." He even suggested, those being the days before the first oil crisis, that much of our protein could be derived from petroleum products.

Perhaps fortunately for our taste buds, none of these have come to pass. But Clarke did predict the communications revolution and an interconnected world thriving on a vast "information grid." Although he thought this revolution would come by way of satellites instead of fiber-optic cables and Internet cafés, the globalized world he envisioned has essentially come true: "The home will have a kind of communications console with a television screen, television camera, computer keyboard, microphone, and probably hard-copy readout." He also foresaw the troubles now facing print journalism, saying that "the newspaper as we know it will be extinct."

Since everyone will be connected remotely, "many people will be able to do most of their work without leaving home—unless their wives insist." (He apparently didn't anticipate the rise of the working woman.) Everyone around the globe would live and work together in step, he said, and time zones will be abolished—or else sleep itself, via some chemical or electronic innovation. The future would give us more free time, and education would be the greatest industry, followed by entertainment. Clarke was hopeful that the world would turn the corner by the 21st century, that "2001 could mark the great divide between barbarism and civilization." Maybe it didn't quite turn out that way, but we can always be optimistic and raise a tall glass of whale milk to the future.

Also in the May 1970 issue, JPL's Ray Newburn (BS '54, MS '55) discussed the prospect of a "Grand Tour" of the outer planets, using a then-untried method of gravitational slingshots to take advantage of a 1976 planetary alignment that had last occurred during the Jefferson administration and wouldn't happen again until 2148. Accompanying the article were two grainy, telescopic photos of Jupiter and Saturn, and neither Newburn nor anyone else could've anticipated the stunning images that the grand tourists—the Voyager spacecraft—would return.

Beyond Saturn, the planets were just fuzzy blobs. Newburn noted that we weren't even sure of Neptune's size, much less its density. Seeing the gas giants up close was a once-in-several-lifetimes opportunity. "We have a chance to do that with reasonable economy this decade; otherwise we must develop new vehicles with greater performance and spacecraft with very long lives, or wait until the middle of the 22nd century." Fortunately, we didn't wait. *—MW* E&S

Caltech

FSC Mixed Sources
Product group from well-managed forests and recycled wood or fiber
Cert no.
www.fsc.org
© 1996 Forest Stewardship Council

Engineering & Science is printed on recycled paper containing a minimum of 10 percent postconsumer waste. The paper's virgin pulp comes from trees grown in the U.S. in accordance with the Forest Stewardship Council's sustainability standards.

VOLUME LXXIII, NUMBER 2, SPRING 2010

e&S Engineering & Science

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 Caltech 1-71, Pasadena, CA 91125. Third-class postage paid at Pasadena, CA. All rights reserved. Reproduction of material contained herein forbidden without authorization. © 2010, California Institute of Technology. Published by Caltech and the Alumni Association. Telephone: 626-395-3630.

Visit us online at **http://EandS.caltech.edu** for videos and web links.

PICTURE CREDITS
Cover, 23 — NASA/JPL; 2, 14, 16, 17, 22, 34, 36 — Bill Youngblood; 3 — Mike Rogers; 5 — Doug Smith; 7 — Caltech Archives, 8 — Vicky Chiu; 12 — Gaby Maimon and Michael Dickinson; 15, 18 — CDMS Collaboration; 17 — NASA/CXC/CfA/M. Markevitch et al., NASA/STScI, Magellan/U.Arizona/D.Clowe et al., NASA/STScI, ESO WFI, Magellan/U.Arizona/D.Clowe et al.; 18, 19 — Fermilab; 22-23, 24, 26–31 — NASA/JPL/SSI; 26, 27 — NASA/JPL/University of Colorado; 30 — NASA/JPL/USGS; 32-33, 37 — Doug Cummings; 33 — Joe Sill

SCHLINGER LAB DEDICATED

Caltech's third new building in as many issues of *E&S* opened its doors on a blustery March 9. Almost a decade in the planning, the Warren and Katharine Schlinger Laboratory for Chemistry and Chemical Engineering is Caltech's first new facility specifically and exclusively designed around the research needs of chemists and chemical engineers since the construction of the adjacent Noyes Laboratory of Chemical Physics in 1967. At the dedication, Warren Schlinger noted that there had been some growth in the years since he arrived on campus as a freshman in 1941, when chemical engineering "had a department made up of two professors and a secretary"—Katharine Stewart, whom he married the year he got his master's degree. (Schlinger contributed to the faculty's growth by staying on as an instructor until 1953.)

Like the recently opened Cahill and Annenberg buildings, the Schlinger Lab is on track to earn a gold certification under the LEED (Leadership in Energy and Environmental Design) Green Building Rating System. Besides using locally derived or recycled building materials, the Schlinger Lab uses 28 percent less energy and 30 percent less water than typical build-

From left: Will Webster (BS '49); Katharine Schlinger; Warren Schlinger (BS '44, MS '46, PhD '49); Chemistry and Chemical Engineering Division Chair Jacqueline Barton, the Hanisch Memorial Professor and professor of chemistry; and President Chameau cut the ribbon.

ings designed for chemical research. Meeting the stringent energy savings required for a gold rating was particularly challenging, as the fume hoods gulp electricity 24/7.

Fume hoods, for the chemically declined, are the enclosed cabinets in which experiments are done. The front side of the hood is a shatterproof windowpane that can be raised for access or lowered until it is almost, but not quite, shut. Powerful fans up on the roof suck a steady draft of room air in under the sash and through the hood in order to keep noxious vapors away from the lab's occupants. The higher the sash is raised, the more air whooshes through the hood and the harder the fans have to work.

Above each of the Schlinger's hoods is an electric eye that constantly scans its vicinity. If nobody is around, the sash automatically lowers to the fully closed position, minimizing the volume of air being pulled in. (A second eye on the sash's underside keeps a lookout for protruding glassware or other objects, stopping the descent if the beam is broken.) This high-tech hood design was pioneered in Europe but is new to the States.

Another European innovation new here can be found in Schlinger's rotary evaporators, which are vacuum-assisted stills. Banks of rotovaps, as they are affectionately known, are essential to any synthetic-chemistry lab—whenever you dissolve something to make it react, you eventually have to get rid of the solvent in order to retrieve your product. A rotovap needs a strong vacuum to get the solvent out as fast as possible, which means either a centralized system with heavy-duty piping, or lots and lots of individual vacuum pumps—noisy, sewing-machine-sized beasts that like to leak their oil all over the lab floor. Instead, each of Schlinger's rotovaps gets its suction from a pump the size of a large paperback book, efficiently and quietly powered by the campus's compressed-air system.

Energy-efficient double-glazed floor-to-ceiling windows flood the labs with natural light, a feat that was made possible by relocating all the plumbing and ducting—normally carried from floor to floor by a "wet" wall—to within a set of elliptical pillars out in the central hallway. By contrast, the entire west facade of Noyes Lab is windowless, hiding the giant utility core that serves the labs. Schlinger's pillars, inlaid with green glass tiles, complement the rich maple accent panels and similarly hued flooring to give an effect reminiscent of the grand corridors of the *Queen Mary*.

Other eco-friendly features include a "bio-swale" on the north side that

Above: Even while you're standing at a fume hood, the outdoors beckons. Below: The Grand Promenade. To see more pictures, check out the slideshow at http://images.caltech.edu/slideshows/Schlinger-architecture/.

collects the runoff from the adjoining planters, sidewalks, and parking stalls, filtering it naturally before returning it to the groundwater. There's also a dedicated room in the basement for collecting and sorting recyclables.

Weather permitting, you can even get in touch with nature without ever leaving the building. One entire wall of the first-floor classroom/conference room folds up into the ceiling like a set of glass garage doors, opening onto a courtyard.

As Division Chair Jacqueline Barton remarked at the dedication, "When you bring chemists and chemical engineers together in one laboratory, the results will be far greater than the sum of the parts." Schlinger's reconfigurable lab spaces will house the research groups of three chemists and three chemical engineers, working in fields ranging from drug design to pollution control. The Center for Catalysis and Chemical Synthesis will also move in, and there's enough room remaining for two new hires.

The building's architects, Bohlin Cywinski Jackson, are known for sustainable design and have done labs and other academic buildings across the nation. Rudolph and Sletten was the general contractor.

Besides the Schlingers, support for the building and its research was provided by the Gordon and Betty Moore Foundation, Will and Helen Webster, Victor and Elizabeth Atkins, the John Stauffer Charitable Trust, Barbara Dickinson, the Ralph M. Parsons Foundation, John Willard Jones (BS '41), Patricia Beckman, and Gregory P. Stone (BS '74, MS '74). — *DS* e&s

OF FLIES AND MEN

The poet and mystic William Blake saw a world in a grain of sand. A fly's brain is scarcely larger, yet Caltech scientists see in it a window for exploring the biological roots of our own behavior and emotions. The brain of *Drosophila melanogaster*, the common fruit fly, contains barely 20,000 neurons—yet two recent papers from the lab of David Anderson, Benzer Professor of Biology and an investigator with the Howard Hughes Medical Institute, offer glimpses into its genetic hardwiring that may throw light on what makes us tick.

For example, both inconsiderate boors and unthinking flies will elbow their way to the front of the press at a crowded lunch counter, causing the less assertive to go elsewhere. Now grad student Liming Wang and Anderson, writing in the January 10 edition of *Nature*, have identified an aggression-promoting pheromone that appears to help drive competitors away from a crowded piece of, say, overripe banana, and pinpointed the neurons in the fly's antennae that detect it.

Pheromones are chemicals used by particular species to communicate with their own kind, but proving that a pheromone released by the insects themselves—rather than being provided in a synthetic form by inquisitive scientists—normally controls aggressive behavior "required the ability to experimentally interfere with the insects' capacity to sense the pheromone," Anderson notes. "And that, in turn, meant identifying the receptor molecules that detect aggression pheromones, and finding the olfactory sensory neurons that express these receptors." According to Wang, the paper's first author, the only insect meeting these requirements was *Drosophila melanogaster*. "The genetic and molecular architecture of *Drosophila*'s olfactory system is well understood," he explains. "One can easily test whether a specific receptor or neuron is involved in a given behavior."

Wang discovered that 11-*cis*-Vaccenyl Acetate (cVA), a pheromone present in the male fly's cuticle, or exoskeleton, promotes aggression in pairs of male flies. An aggressive fly will "lunge," rearing up on its hind legs and snapping its forelegs down on its opponent. When Wang and Anderson added synthetic cVA to the "arena" in which combatant flies were tested, the frequency of lunges dramatically increased. Building on earlier work elsewhere that had identified cVA's receptors, Wang next showed that silencing the cVA-sensitive neurons in the antennae mellowed the flies out.

To find out whether natural cVA from other flies had the same effect, Wang and Anderson then trapped between 20 and 100 "donor" male flies—so called because they donate their pheromones into the surrounding environment—in a tiny cage surrounded by a fine mesh screen. The screen allowed the pheromones to escape, but not the flies. A pair of "tester" males would be placed on top of the cage, where they could sense the pheromone but not interact with the donors. "Remarkably," says

On March 9, Mylar dirigibles battled for the skies—or at least for the airspace within Brown Gymnasium—in "Revenge of the Hindenberg," this year's installment of the ME 72 (Engineering Design Laboratory) contest. For full coverage, see http://weblab.caltech.edu/features/16.

Anderson, "the presence of the caged donor flies strongly increased aggression between the tester flies, and this aggression-promoting effect increased with the number of donors." And again, the testers' testiness was assuaged by inactivating their cVA-sensing neurons.

Which brings us back to the lunch counter—or more aptly, the free food at happy hour. Male flies are attracted to food not only to eat, but also to mate with feeding females. And, of course, the more guys there are, the harder it gets to score. Since feisty flies tend to chase away their competitors, an aggression-promoting pheromone might keep the number of males down to an equitable level.

Wang tested this hypothesis by allowing a small number of flies to compete for a limited food supply, after genetically manipulating their cVA-receptor neurons to make them more excitable. The flies quickly dispersed. "They fought one another until a dominant fly became 'king of the hill' and drove the others away," Anderson explains.

According to Wang and Anderson, this suggests that when the population of male flies reaches a certain density, the concentration of cVA rises to a level that promotes aggression, forcing some of the flies off the food. Their departure decreases the ambient concentration of the pheromone, decreasing aggression. "The population becomes stabilized at an optimal density until more flies become attracted to the food, and the cycle repeats itself," says Wang.

Because pheromones evolved as "private" communications channels within a given species, it's unlikely the fly pheromone would work on us. However, that doesn't necessarily mean that humans lack aggression pheromones, Anderson notes. They've been discovered in mice, which are evolutionarily closer to us than flies, so it's possible we might have our own as well. But whether such pheromones can keep lines short at the buffet, Anderson remarks, "only time will tell."

Anderson's lab has also seen signs of a primitive emotion-like behavior, specifically a state of agitation, that might illuminate the relationship between the neurotransmitter dopamine and attention deficit hyperactivity disorder (ADHD). Most of *Drosophila*'s genes are also found in humans—including those for the neurons that produce dopamine and serotonin, both of which have been implicated in psychiatric disorders.

A team led by then-postdoc Tim Lebestky found that a rapid succession of brief, brisk puffs of air caused flies to run around their test chamber in what Anderson calls a "frantic manner" for several minutes after the last puff. "Even after the flies had calmed down," he adds, "they remained hypersensitive to a single air puff." These "hyperactive" flies were picked out from the crowd via an automated machine-vision-based system developed in the lab of Anderson's colleague Pietro Perona, the Puckett Professor of Electrical Engineering. These flies proved to have a mutation called *DopR* that inactivated a dopamine receptor known as D1—a result that was published in the November 25, 2009, issue of *Neuron*.

This discovery dovetails with what is known about ADHD, which is characterized by impulsivity, hyperactivity, and a short attention span, and is often treated with drugs such as Ritalin that increase dopamine levels in the brain. The way the mutant flies responded to the air puffs is, moreover, "reminiscent of how individuals with ADHD display hypersensitivity to environmental stimuli and are more easily aroused by such influences," says Anderson. Furthermore, ADHD often goes hand in hand with learning disabilities, and Anderson's collaborators at Penn State have shown that flies with the *DopR* mutation can't learn to associate a particular odor with an electric shock. They don't avoid the odor afterward, while flies without the mutation quickly catch on.

It's often assumed that ADHD kids have difficulty learning precisely because they are hyperactive and easily distracted. But this work shows that hyperactivity and learning disabilities are unconnected—in flies, at least. "We could separately 'rescue' the hyperactivity and learning deficits in a completely independent manner," says Anderson, "by genetically restoring the dopamine receptor to different regions of the fly's brain." If it turns out that ADHD works in a similar way, Anderson believes that it may be better to develop drugs to treat the two issues separately. The broad-spectrum pharmaceuticals now used to attack both at once tend to have undesirable side effects.

Besides Lebestky, Anderson, and Perona, the other people involved in the work are Caltech biology research technician Jung-Sook Chang, then-postdocs Heiko Dankert and Lihi Zelnik; Young-Cho Kim and Kyung-An Han from Penn State; and Fred Wolf from UC San Francisco.

That flies exhibit emotion-like behaviors controlled by some of the same brain chemicals as in humans "opens up the possibility of applying the powerful genetics of this model organism to understanding how these chemicals influence behavior through their actions on specific brain circuits," says Anderson. "While the specific details of where and how this occurs are likely to be different in flies and in humans, the basic principles are likely to be evolutionarily conserved, and may aid in our understanding of what goes wrong in disorders such as ADHD."

The research described in both papers was supported by grants from the National Science Foundation and the Howard Hughes Medical Institute.
—*LO* e&s

ANOTHER ALUM GOES TO WASHINGTON

The brain drain from Pasadena to Foggy Bottom continues. Karina Edmonds, MS '93, PhD '98, director of JPL technology transfer, joined the Department of Energy as its first technology transfer coordinator on April 12. She will work with the DOE's national laboratories to help accelerate the process of moving discoveries from the laboratory to the private sector.

Edmonds, an engineer by training, has held key technology-transfer positions at Caltech and JPL for over a decade. As director of JPL technology transfer, she was responsible for technology licensing, managing the JPL patent portfolio, and assisting Caltech/JPL start-up companies.

This is the first time the DOE has appointed a full-time person to fill the role, which was created by the Energy Policy Act of 2005. "I am pleased to have Karina join our team," says Energy Secretary Steven Chu, who was also Caltech's commencement speaker last year. "Having Karina oversee a coordinated, strategic effort on behalf of the department will help increase the rate of successful technology transfers, creating clean-energy jobs and providing more solutions to our energy challenges." —*AB* e&s

The six figures on Calder's arches represent (from left) Nature, Art, Energy, Science, Imagination, and Law.

CALTECH TURNS 100—AGAIN

Although founded in 1891, Caltech can once again celebrate its 100th anniversary. On June 8, 1910, the first building on the present-day campus was dedicated before local dignitaries and a large public audience. Dubbed Pasadena Hall (and renamed Throop Hall in 1920), it was hailed as a monument to civic pride. The first students to be educated in it arrived the following September—30 in total, all male, and all enrolled in a college-level engineering curriculum. Tuition was $150 annually.

Throop Polytechnic Institute, as it was then known, had just split apart at the seams. The old Throop had evolved into an agglomeration of six schools, teaching at levels from elementary to collegiate, with a heavy emphasis in the upper division on such practical skills as stenography, typing, and operating machine tools.

The leap from vo-tech to high-tech was the work of noted solar astronomer George Ellery Hale, who had come west in 1903 to be the founding director of the Mount Wilson Observatory. Hale was soon deeply in the flow of Pasadena's civic, cultural, and educational schemes, becoming a tireless booster of Southern California in general and Pasadena in particular. He soon became bent on establishing a local technical school to train engineers (construed at that time to mean men only) to meet the needs of a booming region—in particular, to bring water and electricity over the mountains to a sun-drenched but utility-starved Los Angeles basin. By 1907 he had begun a campaign for the creation of a "high-grade institute of technology" in Pasadena and was elected to the Throop board of trustees. In that same year, an anonymous benefactor secured a site for a new, expanded campus—some 22 acres of orange groves dotted with stately oak trees in the southeast part of the city.

Throop's original campus—acquired after a start-up year in the old Wooster Block, still a presence on the corner of Fair Oaks Avenue and Green Street in the heart of Old Pasadena—crammed all six schools into three buildings at Lincoln Avenue and Fair Oaks, a site that is today under the 210 freeway. Hale envisioned the new campus two miles east as an opportunity for an idealized building scheme in harmony with a new civic center, a campus whose laboratories would be fitted out with the latest and best equipment. Such an institute would redound to the glory of Pasadena and would surely inspire the generosity of Pasadena's well-to-do residents.

Hale was right. The mission-style structure by architects Myron Hunt and Elmer Grey was paid for entirely by local subscriptions, to the tune of approximately $165,000. The arcaded entrance was adorned by a set of reliefs created by Pasadena's Alexander Stirling Calder, whose son would invent the mobile. Touted at the time as the most significant artwork

Throop Hall (with Dabney Hall to the left) in April 1965. The president, the provost, the treasurer, and the deans had offices on the first floor. Various business offices—payroll, personnel, accounting, central files, and so on—occupied the second floor. Ed Hutchings, editor of *E&S*, lived in the basement with the news bureau, the alumni office, and most of development.

Throop Hall was demolished after the San Fernando earthquake, and the Calder Arches now adorn the facade of the Arnold and Mabel Beckman Laboratory of Chemical Synthesis. The Throop site is now a vest-pocket park in the middle of campus—a perfect place for a photo op. Here Lemelson winner Heather Agnew (right) and finalist Yvonne Chen enjoy their accolades.

in the city, the elaborate, allegorical figures were, in Calder's words, "to give plastic utterance to the aims and scope of the school."

The new building's 62 rooms housed what the 1910 catalog boasted as being "the only college devoted primarily to Technology west of the Mississippi River." Meanwhile, the other five schools were closed down or divested. The elementary school moved to a new location a block west of the new campus and became the Polytechnic School. Throop Academy remained at the old campus and eventually merged with a new public high school. And, after almost becoming UC Pasadena in 1911 and completing the Gates Laboratory of Chemistry in 1917, Throop College of Technology rebranded itself as the California Institute of Technology in 1920—so we have another 10 years to wait for that party.

An online exhibit about the 1910 campus may be found at the Caltech Archives website: http://archives.caltech.edu/. —*SE* E&S

LEMELSON WINNERS ANNOUNCED

Deep in the Amazon, a woman is keeling over with stomach pains and vomiting. Does she just have the flu, or is she one of two billion people worldwide who has been afflicted with Hepatitis B, a potentially deadly liver disease? Today's diagnostic tools are too delicate for health workers to use in the steamy environment of a remote jungle. But in the future, a drop of blood from a prick of the finger and a cheap, simple device that works in nearly all conditions may change that. Heather Agnew (PhD '10) and Jim Heath, the Gilloon Professor and professor of chemistry, are working to make such devices a reality. For her role in this effort, Agnew has won the $30,000 2010 Lemelson-MIT Caltech Student Prize.

A diagnostic test, or assay, can measure the amount of a protein specific to some disease by allowing it to bind to another molecule, called an antibody, that is tailor-made to recognize it. Assays can be packaged into easy-to-use kits for diagnosis outside the lab, and they're not restricted to blood. For example, the home pregnancy test assays a hormone called human chorionic gonadotropin in urine.

The problem with such assays, though, is that the antibodies themselves are proteins, sensitive to heat, humidity, and other factors. For instance, HIV assays have to be performed within hours of opening the package or the antibodies degrade, Agnew says. But the developing world, which needs such simple diagnostic tools the most, isn't always air-conditioned. Furthermore, antibodies are expensive to produce. Today, many tests look for just one or two proteins, Agnew says. But diseases like cancer are complex, so an accurate diagnosis might require the measurement of more than a dozen proteins, each by its own antibody.

Agnew and her colleagues are building cheap, durable antibody replacements called protein-capture agents out of synthetic peptides, which are relatively short chains of amino acids—the building blocks of proteins. Peptides are cheap to make, and can be designed to have all sorts of nice properties, including heat resistance and biological or chemical stability. But since they're small molecules, they don't stick to their target proteins as well as antibodies do.

Reasoning that two peptides of middling stickiness might do the trick if they worked together, Agnew and her coworkers tested millions of them. And here the project got help from the target protein itself—when appropriately primed versions of the peptides recognized the protein and bound to it, it held them in just the right orientations that they could "click" together to create a new molecule that is 10 to 100 times better at binding to the target protein than either peptide alone. Repeating the process to add a third peptide further enhances the binding.

As for durability, Heath's benchmark is what he calls the Pasadena Test: will it work even after a year spent baking in the trunk of his car? Agnew says her protein-capture agents have withstood airplane travel and years of sitting on a shelf in her office.

A second award of $10,000 went to Yvonne Chen (MS '07), a grad student working with Christina Smolke, a former assistant professor of chemical engineering at Caltech who's now at Stanford. Chen developed a way to help T cells fight cancer. T cells are a part of the body's protective army, and other researchers have been able to engineer them to attack cancerous tumors. "We can keep putting them in the blood supply until they home in on the tumor," Chen explains. "The problem is that they die really quickly." Because T cells are a part of the body's natural immune response, they die by default if they aren't instructed to attack. "Our challenge then is to figure out how to engineer this T-cell population to be sustainable so they can finish killing the tumor cells."

T cells are kept alive by molecules called cytokines. But you can't just inject cytokines into someone to keep the T cells going—you'd need a lot, and too much would put the patient into shock. One solution is to engineer the T cell to produce its own cytokines. But you also have to regulate cytokine production carefully, because an excess will cause the T cells to reproduce nonstop, resulting in leukemia.

With Smolke, Michael Jensen from City of Hope medical center, and other researchers, Chen made a molecule of RNA—which is similar to DNA—that acts like a switch, turning cytokine production on when exposed to theophylline, a caffeine-like molecule (see *E&S* 2005, No. 4). When the theophylline infusion stops, so does cytokine production, and the T cell dies. This is just a demo, as large doses of theophylline can cause an irregular heartbeat and even death. Fortunately, the RNA switch can easily be designed so that it responds to a harmless molecule, such as a vitamin. Chen is now working to make it more versatile and easier to control.

The Lemelson-MIT Caltech Prize is funded by the Lemelson-MIT Program, founded in 1994 by Jerome H. Lemelson to inspire young innovators. Chen's prize as a finalist was donated by Michael Hunkapiller (PhD '74). Lemelson-MIT student prizes are also at MIT, Rensselaer Polytechnic Institute, and the University of Illinois at Urbana Champaign. The Caltech prize was first awarded last year.

—*MW* e&s

VOYAGER NEARS THE LOCAL FLUFF

Our solar system is plunging through a vast cloud of wispy gas called the local interstellar cloud, also known as the "Local Fluff." About 30 light-years wide, the Fluff is made of 6,000°C hydrogen and helium. The Fluff is about twice as dense as the interstellar meduim surrounding it, and what holds it together has been a mystery—until now, thanks to a discovery by JPL's twin Voyager spacecraft.

By rights, the Fluff shouldn't exist. A group of nearby stars exploded about 10 million years ago, and the resulting blast of million-degree gas is now blowing past us. The Fluff is neither hot enough nor dense enough to withstand the onslaught, says Merav Opher, a former JPL postdoc now on the faculty at George Mason University. But in the December 24, 2009, issue of *Nature*, Opher and her colleagues reported that the latest data from Voyager 2 reveal a magnetic field strong enough to enable the Fluff to push back. "Voyager data show that the Fluff is much more strongly magnetized than anyone had previously suspected—between four and five microgauss," Opher told Science@NASA. "This magnetic field can provide the extra pressure required to resist destruction."

Previous estimates of the Fluff's field had been in the 1.8 to 2.5 microgauss range. By comparison, Earth's magnetic field is about half a gauss, or roughly a million times stronger.

Inside the Fluff—and encompassing us—is a 10-billion-kilometer-wide bubble called the heliosphere, which helps shield us from constant bombardment by high-energy cosmic rays from the depths of space. The heliosphere is kept inflated by the solar wind, a stream of charged particles emitted by the sun, so its size is determined by the balance of forces between the solar wind pushing out and the local interstellar cloud pressing back. In 2004 and

2007, respectively, Voyagers 1 and 2 crossed into the heliosphere's outer layers, a region called the heliosheath. (See *E&S* 2008, No. 3.) Once there, they could measure the size of the heliosphere, allowing scientists to calculate how much pressure the Fluff is exerting on it. This pressure, in turn, partly depends on the strength of the Fluff's magnetic field.

This discovery raises the possibility that other clouds in our galactic neighborhood are also strongly magnetized, and when the solar system collides with them, they will push back even harder. If the heliosphere is further compressed, more cosmic rays might reach Earth. "There could be interesting times ahead," Opher says. But there's no need to get out the tinfoil hats quite yet—we won't run into the next cloud for hundreds of thousands of years. —*MW* e&s

TUNE IN TO "TODAY"

For more on what's happening with Caltech and the Caltech community, check out the articles on Caltech Today, which offers online coverage of the Institute and its activities. There you'll find press releases covering the latest Institute research and feature articles highlighting faculty, student, alumni, and campus activities. Click on the links with icons for the latest stories or scroll down to the Feature and News archives at the foot of the page and click on those links to browse through older articles.

Caltech Today can be accessed through its link on the Caltech home page or directly at http://today.caltech.edu/. You can also subscribe via email or RSS feed. e&s

MONEY ON YOUR MIND

It's perhaps not surprising that aversion to losing money is hardwired into our brains, but a sense of fairness seems to be as well. These are just two results from recent work at Caltech's Brain Imaging Center, where a multidisciplinary team of biologists and social scientists are using functional Magnetic Resonance Imaging (fMRI) to map behavior onto brain structure with millimetric precision.

An fMRI scanner tracks blood flow in the brain as a proxy for brain activity. The test subject lies in the scanner and is then asked a series of questions or told to perform some other sort of mental activity, such as memorizing a list of names, while the experimenters literally watch him or her think. Many experiments use pairs of volunteers, each in their own scanner, trying to outwit each other in various strategy games where cash is on the line.

It turns out that the fear of losing money lives in the amygdalae, two almond-shaped clusters of tissue located in the medial temporal lobes. (The amygdala registers rapid emotional reactions and appears to play a role in depression, anxiety, and autism.) Benedetto de Martino, a visiting researcher from University College London; Colin Camerer, the Kirby Professor of Behavioral Economics; and Ralph Adolphs (PhD '92), the Bren Professor of Psychology and Neuroscience and professor of biology, found the seat of this fear by studying two patients whose amygdalae had been destroyed by a very rare genetic disease.

These two people, as well as other volunteers, were each given $50 in cash and then offered a series of bets on the outcome of a computer-generated coin toss. Each potential wager had the same odds, 50/50, but a different ratio of payout to loss. For example, you might get the chance to win $20 or lose $5 (a risk most people will accept), or you might stand to lose $20 for the same $20 return (a bet most people will decline). In general, people shied away from the prospect of large losses, so even the proposition of winning $20 versus losing $15 got few takers, "even though the net expected outcome is positive," Adolphs says.

Neither of the amygdala-damaged patients were fazed by the prospect of losing money, taking risky gambles much more often than control subjects. "We think this shows that the amygdala is critical for triggering a sense of caution," explains Camerer. This function, he says, may be similar to the amygdala's role in fear and anxiety. "Loss aversion has been observed in many economics studies, from monkeys trading tokens for food to people on high-stakes game shows," he adds, "but this is the first clear evidence of a special brain structure that is responsible for fear of such losses."

A paper on this research appeared in the February 23 issue of the *Proceedings of the National Academy of Sciences*. The work was supported by the Gordon and Betty Moore Foundation, the Human Frontier Science Program, the Wellcome Trust, the National Institutes of Health, the Simons Foundation, and a Global Center of Excellence grant from the Japanese government.

Another study, by Professor of Psychology John O'Doherty, Camerer, then-postdoc Elizabeth Tricomi, and Associate Professor of Economics Antonio Rangel (BS '93), looked at the brain's reward centers. It's long been known that we don't like inequality, especially when it comes to money. Tell two people working the same job that their salaries are different, and there's going to be trouble, notes O'Doherty. "It's not just the application of a social rule or convention; there's really something about the basic processing of rewards in the brain that reflects these considerations."

The experimenters watched how the ventromedial prefrontal cortex (VMPFC) and the ventral striatum—two well-known reward centers in the brain—reacted to the prospect of being offered various amounts of money. But there was a twist—the 40 volunteers were paired off beforehand, and one person in each pair was given an extra $50 before the experiments even began. Then, in each trial, the pair would be told how much more money they could potentially get—from zero dollars up to another $50—in a payout scheme selected at random at the end of the run.

As it turned out, the way the volunteers—or, to be more precise, the reward centers in their brains—reacted depended strongly upon whether the volunteer was the "poor" or the "rich" member of the pair. "People who started out poor had a strong reaction to getting money, and essentially no reaction to money going to another person," Camerer says. "By itself, that wasn't too surprising." What was surprising was the other side of the coin—"people who started out rich had a stronger reaction to other people getting money than to themselves getting money. In other words, their brains liked it better when their poorer partner got the money."

"We now know that these areas are not just self-interested," adds O'Doherty. "They don't exclusively respond to the rewards that one gets as an individual." Instead, contrary to the prevailing wisdom about human nature, the brain evaluates the overall equity of the situation. "It shows that the basic reward structures in the human brain are sensitive to even subtle differences in social context."

Camerer, too, found the results thought provoking. "We economists have a widespread view that most people are basically self-interested, and won't try to help other people," he says. "But if that were true, you wouldn't see these sort of reactions to other people getting money." Still, he says, the rich may have been at least partly motivated by self-interest—or a reduction of their own discomfort. "We think that, for the people who start out rich, seeing another person get money reduces their guilt over having more than the others."

O'Doherty says that the next step is to attempt to figure out how these reactions translate into changes in behavior. "For example, the person who finds out they're being paid less than someone else for doing the same job might end up working less hard. It will be interesting to try to understand the brain mechanisms that underlie such changes."

These findings were published in the February 25 issue of *Nature*. The project was funded by grants from the National Science Foundation, the Human Frontier Science Program, the Gordon and Betty Moore Foundation, and the Caltech Brain Imaging Center. —*KS/LO* E&S

"People who started out rich had a stronger reaction . . . . Their brains liked it better when their poorer partner got the money."

GETTING INSIDE A FLY'S HEAD

What goes on in the tiny brain of a fruit fly? We're beginning to find out, now that Michael Dickinson, the Zarem Professor of Bioengineering, and postdocs Gaby Maimon and Andrew Straw have succeeded in recording the activity of individual brain cells as the fly flies. This is no mean feat, considering that each fly is only about 2.5 millimeters long.

"Researchers have recorded the neural-cell activity of fruit flies before, but only in animals that had been stuck or glued down," Dickinson explains. "Gaby was able to develop a preparation where the animal is tethered"—its head clamped into place—"but free to flap its wings." By slicing off a patch of the hard cuticle covering the brain, "we were able to target our electrodes onto genetically marked neurons," he says. As the electrodes took data, high-speed digital cameras simultaneously recorded the flies' behavior.

The study focused on a set of visual-system neurons that "basically help the fly detect when its body posture changes" in order to maintain stable flight, Dickinson says. When the wings started flapping, these cells immediately ramped up their activity. "The neurons' responses to visual motion roughly double when the flies begin to fly, which suggests that the system is more sensitive during flight," Dickinson says. "The increase is very abrupt. It's not at all a subtle change, and so we suspect that there is a neurochemical quickly released during flight that sets the animal's brain in this different state."

Previous studies in locusts—which are far bigger and thus far easier to study—had suggested the existence of this effect. However, the genetics of locusts are not nearly as well understood as those of *Drosophila*. Now, says Dickinson, it should be possible to "figure out specifically what causes the change in sensitivity. Is the system turned off when the fly is on the ground? What neurochemicals are involved? We can use all the genetic tricks that are available in fruit flies to get a better idea of what is going on." Adds Maimon, "Sensory neurons in many species—including birds, rodents, and primates—change their response strength depending on the behavioral state of the animal, but why these changes take place is not entirely clear."

The researchers also plan to spy on olfactory and motor cells to see if they display similar behavior. "The question is, 'Is the entire brain completely different in flight?'" Dickinson says. "We suspect that this phenomenon is not unique to the visual cells we have studied. Most cells care whether the animal is flying or not."

A paper describing the research was published in the March issue of *Nature Neuroscience*; the work was funded by the National Science Foundation and a Caltech Della Martin Fellowship. —*KS*

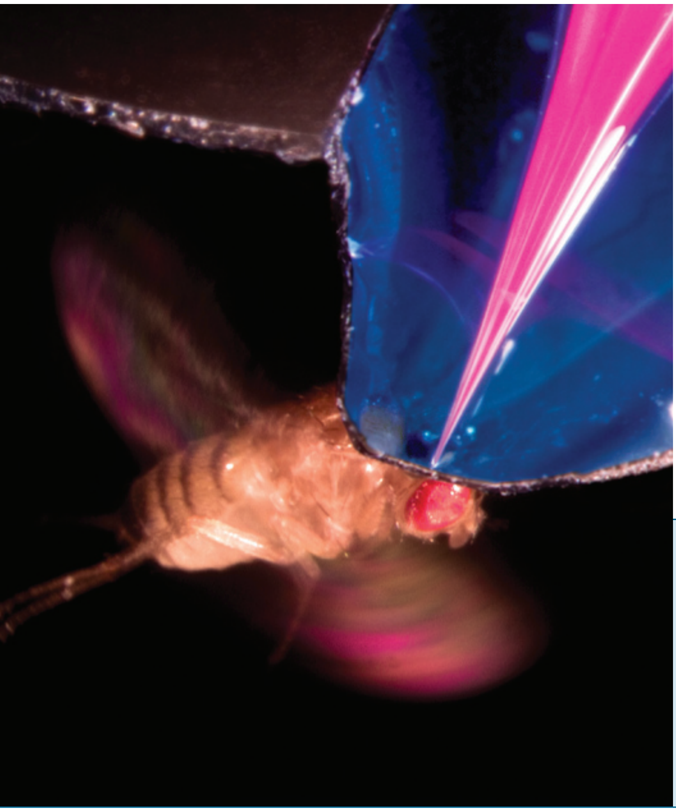

This fruit fly has a dye-filled glass electrode (pink) inserted into its brain. The fly's head is clamped to the underside of a reservoir filled with a sterile saline solution (colored blue here) that bathes the electrode and the brain. At rest, the fly clings to the reservoir; a gentle puff of air starts it flapping its wings in tethered flight.

OLD MAGAZINES NEVER DIE . . .

The Intel Science Talent Search, formerly the Westinghouse Science Talent Search, is to high-school science fairs what the World Series is to sandlot baseball. The grand prize is $100,000, and recent winners of this nationwide competition have done such things as creating a 50-gene model for predicting the probability of a specific colon cancer recurring, building a Littrow-type spectrograph, and designing a nanosensor for neurotoxins.

This year's top honor went to Erika DeBenedictis of Albuquerque, New Mexico, for "a software navigation system that would allow spacecraft to exploit low-energy orbits . . . for more efficient transit routes through the solar system."

DeBenedictis built on research by JPL's Martin Lo (BS '75 and a Science Talent Search winner himself), in collaboration with control and dynamical systems professor Jerrold Marsden's research group, on what Lo calls the "Interplanetary Superhighway"—a set of low-energy routes connecting every massive body in the solar system through the intersections of rotating Poincaré manifolds. In fact, Lo, a colleague of Erika's father, Sandia National Lab's Erik DeBenedicits (BS '78, PhD '83), helped her get started on a precursor project in 2007–8.

If the Interplanetary Superhighway sounds vaguely familiar, it's because an article on it appeared in *E&S* in 2002, when DeBenedicits would have been a fifth-grader. In a presentation she gave at JPL on April 15, she cited *E&S* as her inspiration.

Contacted by email, she elaborated, "I think what happened (as with most interesting science articles) was that I saw something I liked and asked my dad to explain it to me. That's why when I thought of it a few years later he remembered it too and was able to find it again.

"You would probably be surprised how much difference the articles you write make—*E&S* is one of my favorite magazines to flip through and look at the cool stuff."

DeBenedictis will be matriculating at Caltech in the fall, and hopes to work at JPL when she graduates.
—*DS* e&s

MAKING BOOK ON THE LHC

As you no doubt know by now, the Large Hadron Collider, or LHC, is back up and running again at a stable, record-setting collision energy of seven trillion electron volts. The LHC was switched on with great fanfare in September 2008 (see "Beam On!," *E&S* 2008, No. 3) and shut down again nine days later due to a faulty electrical connection that led to a massive coolant leak and ultimately damaged 53 of the more than 1,600 superconducting magnets. It took over a year to repair everything, and the LHC was restarted again in November 2009, just in time for the regularly scheduled winter shutdown.

Caltech physics faculty, staff, and students pulled an all-nighter to watch the restart on a live video feed from Geneva, where the European Organization for Nuclear Research (CERN) and the LHC are located. And it was a long night—after two false starts, the countercirculating proton beams were finally brought into collision just after lunch in Switzerland, which unfortunately translated into 3:58 a.m. our time.

Meanwhile, according to a press release received by this office, a publicly traded Irish online betting firm named Paddy Power is laying odds on what the LHC will discover first. "The mysterious and previously undetectable form of matter known as Dark Matter is the red-hot 11/10 favourite, followed by Black Holes at 8/1 and Dark Energy at 12/1. God remains the 100/1 outsider."
—*DS* e&s

A Mine for Dark Matter

Researchers have built an experiment 230 stories underground in search of an elusive particle that may be dark matter, the mysterious stuff that makes up nearly a quarter of the universe.

For two weeks in mid-December 2009, the physics world was buzzing with anticipation and speculation. A team of researchers was rumored to have made an astounding discovery—they'd detected dark matter, the unknown stuff that makes up nearly a quarter of the universe. The world's leading experiment to find dark matter, the Cryogenic Dark Matter Search (CDMS), had just finished analyzing its final data set. Word had somehow reached the blogosphere that the results would be published in *Nature*; an announcement to be made in such a widely read and prestigious journal must mean big news. This was soon debunked, but it was too late. The rumor had already drawn the attention of physics blogs, including Cosmic Variance (in a post by Caltech senior research associate Sean Carroll), and those of *New Scientist* and *symmetry* magazines.

To further fan the flames, a pair of talks announcing the results were scheduled

Assistant Professor of Physics Sunil Golwala stands in front of a three-dimensional map of dark matter. Measurements of how light bends around massive galaxy clusters, an effect called gravitational lensing, allow astronomers to better estimate not only how much dark matter is out there, but also how it has governed cosmic evolution. This map, which has time as its third dimension, goes halfway back to the Big Bang (over Golwala's shoulder). As your eye moves to the left, the dark matter goes from smooth to clumpy, dragging the visible matter with it to form the galactic structures of today. This map, the first of its kind (see *E&S* 2007, No. 1), was published in a 2007 paper by then-postdoc Richard Massey and others, including JPL scientist Jason Rhodes and Steele Family Professor of Astronomy Richard Ellis. The data came from the Hubble Space Telescope's Cosmic Evolution Survey, led by Moseley Professor of Astronomy Nick Scoville, which spent nearly 1,000 hours looking at a patch of sky about the size of nine full moons.

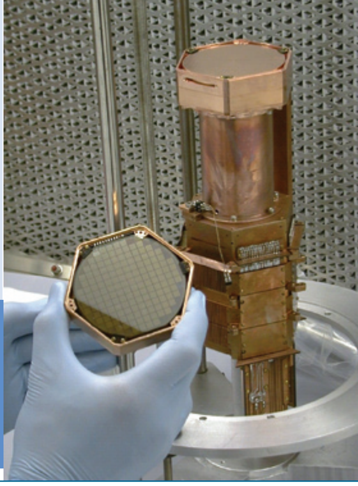

Each CDMS detector is a 230-gram germanium crystal. Six detectors are stacked to form one of the five towers that make up the whole apparatus.

By Marcus Y. Woo

to be given simultaneously at Fermilab in Illinois and the Stanford Linear Accelerator Center (SLAC) in California. These talks were set weeks before, says Assistant Professor of Physics Sunil Golwala, a member of the CDMS team. But, coupled with the rumors, the scheduled talks only added to the rampant speculation. “Then it got crazy for a couple of weeks,” Golwala recalls. The CDMS team had decided early on not to discuss the results before the presentations, to ensure that data wouldn’t be released before a thorough vetting. The secrecy, however, just got people more suspicious. “People came up to me and tried to read my facial expression,” says Jeff Filippini, a postdoc and CDMS team member.

No one knows for sure what dark matter is made of, but so far the best guess is that it consists of a type of particle called a weakly interacting massive particle (WIMP). If physicists and astronomers are right about this, then WIMPs should be all around us, zooming about at hundreds of kilometers per second. But because they hardly interact with regular matter, you can’t see or feel them. There could be billions of them streaming through your body right now. Once in a while, though, a WIMP could crash into an atomic nucleus like a cue ball hitting an eight ball, and that’s the idea behind most dark-matter searches, including CDMS.

The detector consists of 30 hockey-puck-sized crystals of germanium waiting for a WIMP to come along. To block cosmic rays that might confuse the signal, CDMS sits about 230 stories deep in the Soudan Underground Laboratory, a research facility built by the University of Minnesota, Fermilab, and the Minnesota Department of Natural Resources. Why the DNR? Because the lab sits in the bowels of an old iron mine nestled among the lakes and forests at the northeast tip of Minnesota. The CDMS team numbers nearly 80 people from 16 institutions around the world, including Caltech.

Although CDMS is far from alone in trying to catch WIMPs, it’s been the standard-bearer for the past few years. No experiment has yet detected anything, but each silent result narrows down what WIMPs might look like—any theory that predicts something the experiments don’t see has to be refined or ruled out. CDMS has provided the tightest constraints yet, and these latest results, taken over a period of more than a year, have doubled the collaboration’s data. If physicists were close to finding WIMP collisions, then CDMS would have been the first experiment to do so—which explains why people became so anxious upon hearing the rumors. The hype underscores just how momentous a dark-matter discovery would be. “It’s a really exciting topic,” says Golwala, whose two graduate students, Zeeshan Ahmed (MS ’08) and David Moore, did a lot of the number crunching for the new data. “Suppose you have conclusive evidence that you just discovered the dark matter in the universe,” he says. “I mean, that’s just amazing.”

WHAT’S THE MATTER?

For decades, dark matter remained an abstraction, living within the confines of conjecture and theory. Caltech’s Fritz Zwicky—the eccentric, cantankerous iconoclast who was a professor of physics from 1941 to 1968—coined the term nearly 80 years ago. In 1933, he found that galaxies in a group called the Coma Cluster were zipping around a common center of gravity much faster than they should’ve been—at those speeds they should have been flying apart. The only way the galaxies could stay clustered was if there was more mass to them than met the eye—some new type of matter that only interacted with stars, dust, and gas through mutual gravitation. This stuff couldn’t be seen, and was therefore “dark.” Other astronomers, however, didn’t take his pronouncement seriously. (In fact, much of Zwicky’s research was ahead of its time, and now many consider him to be an overlooked genius.)

Not until the 1970s, when Vera Rubin of the Carnegie Institution of Washington measured how fast spiral galaxies spin, did the notion of dark matter gain greater acceptance. The principle behind her discovery is the same as Zwicky’s—she looked at dozens of spiral galaxies and found that the outer stars were circling so fast that the galaxies should’ve been ripping apart. Since then, even more accurate measurements of galaxies and galaxy clusters have revealed a universe filled with mass we can’t see, holding galaxies together like cosmic glue. Nearly all galaxies appear to be embedded

In addition to CDMS, Golwala's research group works on a variety of other topics in observational cosmology, exploring the origin of the universe and the nature of dark matter. Golwala uses Bolocam—a camera built with Andrew Lange, JPL scientists Jamie Bock and Hien Nguyen, and Jason Glenn of the University of Colorado—at the Caltech Submillimeter Observatory to study tiny fluctuations in the cosmic microwave background (CMB) caused by galaxy clusters. Golwala is also developing a new camera called MUSIC with Professor of Physics Jonas Zmuidzinas (BS '81), Nguyen, Glenn, JPL postdoc Jack Sayers (MS '04, PhD '08), and JPL scientists Peter Day (PhD '93) and Rick LeDuc.

Golwala also plays a role in two more experiments that analyze the CMB. Spider, on which postdoc Jeff Filippini is a team member, is a balloon experiment that will fly from Antarctica and observe at millimeter wavelengths. BICEP2, on which postdoc Walt Ogburn is a team member, is a sister experiment situated at the South Pole.

in huge dark-matter halos several times the size of galaxies themselves. According to the latest estimates, about 85 percent of all the matter in the universe is dark. (And it turns out that most of the cosmos isn't even matter—three-fourths is dark energy, an entirely different beast altogether and an even bigger mystery.)

Physicists and astronomers have come up with no end of ideas to account for the universe's invisible mass. Some astronomers have even proposed that our theory of gravity is incomplete, that it behaves differently on cosmic scales. But in 2006, observations of the Bullet Cluster, published by a team of astronomers that included lead author Douglas Clowe (BS '93), Anthony Gonzales (BS '95), and Dennis Zaritsky (BS '86), made a convincing case that dark matter is indeed real. (Clowe, Gonzales, and Stephen Murray [PhD '71] were also among the authors of a lead-up study on the Bullet Cluster's dark matter in 2004.)

Most of the visible matter in galaxy clusters consists of X-ray-emitting hot gas, and the Bullet is actually two galaxy clusters in the midst of a high-speed crash. When two clusters smack into each other, the gas collides, and the clusters slow each other down. Dark matter, on the other hand, hardly interacts with anything, so the two giant blobs of dark matter would pass through each other like ghosts. When astronomers measured how light bends around the Bullet Cluster, an effect called gravitational lensing, they found that most of the mass is not in the hot ball of colliding gas, but in the places where the dark matter would be. For the first time, astronomers had isolated dark matter from visible matter.

Some hypotheses say that dark matter is composed of familiar-but-dim objects, like black holes or brown dwarfs—Jupiter-sized balls of gas too small to form stars. These things have been dubbed massive compact halo objects (MACHOs—since they're obviously not WIMPs), but there doesn't seem to be enough of them out there. Two lines of evidence from the Big Bang have now convinced most astronomers that dark matter isn't normal stuff, made from protons and neutrons, but something completely different.

Right after the Big Bang, the cosmos was a soup of hot plasma. As the universe expanded and cooled from 10^{32} degrees to a balmy 10^{9} degrees, protons and neutrons formed. In this fiery cosmic cauldron, nuclear fusion took place as protons and neutrons slammed into one another, forging the first elements—hydrogen, helium, and lithium—and some of their isotopes, which have varying numbers of neutrons. For example, the nucleus of ordinary hydrogen is a bare proton. Deuterium, the next heaviest hydrogen isotope, has a proton and a neutron, and tritium has a proton and two neutrons. The total density of protons and neutrons dictated the isotopes produced and in what proportions. This primordial process is the only way to make deuterium, so by measuring the abundance of deuterium now, we can gauge the total density of protons and neutrons then. It turns out that there aren't enough protons and neutrons—collectively called baryons—to account for all the mass in the universe.

The cosmic microwave background (CMB), the sky-filling afterglow of the Big Bang, also suggests that dark matter must be different—or nonbaryonic, in physics nomenclature. Observations of the CMB—including those from BOOMERanG, the Antarctic balloon telescope experiment led by the late Goldberger Professor of Physics Andrew Lange—show a stipple that betrays the universe's baryonic density. (See *E&S* 2000, No. 3). "You're looking back in time to a plasma a few hundred thousand years after the Big Bang," Filippini explains. "You can think of it as a snapshot of the way this plasma is frothing and sloshing from place to place." The normal stuff—baryons—is being pushed around by the energy in this cosmic soup, adds Walt Ogburn (BS '99), a postdoc who was also on the CDMS team. "If you have some kind of nonbaryonic dark matter that's not interacting with the protons, then it does its own thing and interacts gravitationally with the soup." So how much the plasma sloshes—its amplitude, which is represented by the prominence of the spots

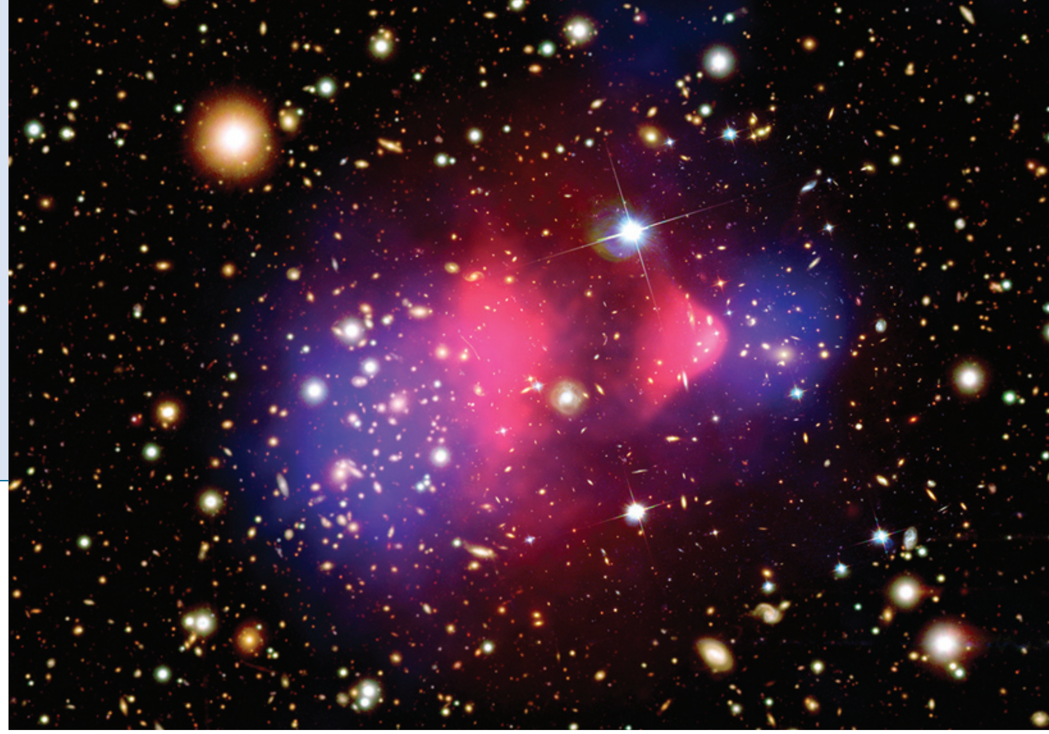

A composite image of the Bullet Cluster, which is two galaxy clusters that began colliding about 3.5 billion years ago. The red represents each cluster's X-ray-emitting gas, which was slowed down by the collision and remains near the crash site. The blue represents the dark matter. Since dark matter hardly interacts with anything, it sailed right on through the other cluster unimpeded, and it's still going.

in the CMB—tells astronomers what kinds of ingredients are in the soup . . . and a lot seems to be that strange, nonbaryonic dark matter.

Theorists have concocted a smorgasbord of potential dark-matter candidates, exotic particles with names like Q-balls, cryptons, and WIMPzillas. But so far, the particle that is considered to be the most credible candidate is the WIMP. "In my mind, it's the best option," says Golwala. A WIMP is about a hundred times more massive than a proton and interacts with other particles through gravity and the weak force (which are two of the four fundamental forces of nature—the others being the electromagnetic force, which is actually the same as the weak force at sufficiently higher energies, and the strong force, which holds nuclei together). What distinguishes WIMPs as good candidates is that they're really not special at all, but a type of particle predicted by many of the latest theories in particle physics.

The Standard Model, which explains how all the fundamental particles interact with one another, is one of the most successful theories in physics. But it's incomplete—for example, it doesn't explain gravity. It also turns out that if you were to solve the relevant equations, you would find that every particle—except the photon or gluon, another massless particle—should have a mass of 10^{19} gigaelectron volts, the so-called Planck mass. Instead, a proton has a mass of about 100 gigaelectron volts—and the Standard Model can't account for the enormous gap.

To fix the problem, physicists have cooked up several theories. One of these is supersymmetry, in which every particle has a new partner—for instance, a quark's "superpartner" is a squark. Another group of theories posits that there are more spatial dimensions than the three we normally experience—up, down; left, right; forward, backward. The extra dimensions are curled up into sizes too small for us to see. The theories of supersymmetry and extra dimensions each suggest that there could be particles with the general properties of WIMPs.

If there are WIMPs, they would've been produced in the early universe, back when it was so dense that despite their weak interactions, they would've been dashing around, smashing into each other. When they did, they would have annihilated one another, bursting into other particles and energy. But as the universe expanded and became less dense, it became harder for WIMPs to find and annihilate each other, and they soon were no longer able to, leaving a bunch of them hanging around with nowhere to go. Calculations predict that the quantity of leftover WIMPs is roughly the same as the estimated amount of dark matter. "It's just a complete coincidence—and that is what got people really excited," Golwala says. "There's no reason this coincidence had to happen, and therefore we should take it seriously." Whereas many other potential dark-matter particles are pure invention, WIMPs are born naturally out of particle physics. "It doesn't rely on one specific theory of how things work—that's the most compelling thing about it."

How exactly a WIMP behaves—how frequently it would smash into the nuclei in CDMS's detectors, for instance—depends on the specific variation on supersymmetry or extra-dimension theory, and there are many. "For every theorist who's thought about it, there's going to be a slightly different theory," Filippini says. Still, most of these ideas require, or are consistent with, WIMPs. "A particle physicist and a cosmologist would both think this kind of particle should exist," he remarks. "That's such a cool coincidence that it makes people actually want to look for these things."

SUCH A WIMP

The Soudan iron mine opened in 1884 and would become Minnesota's deepest, as well as oldest, mine, helping the state become the nation's leader in producing iron

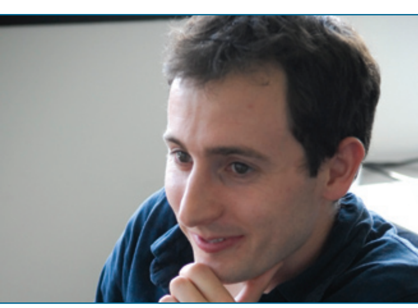

Far left: David Moore, Golwala, Jeff Filippini, and Zeeshan Ahmed in front of the Cahill Center for Astronomy and Astrophysics.

ore. It closed in 1962, and is now a state park as well as a physics lab, home also to a neutrino-detection experiment called MINOS. Most flights take the researchers to Minneapolis, where they then drive for four hours to Soudan—a journey made perilous by Minnesota winters. "California drivers don't do well in the snow," says graduate student Moore, who, along with his fellow grad Ahmed, has spent plenty of time in the mine. "Sometimes with the wind chill, it's minus 40 degrees—it's kind of crazy," adds Ahmed, who played a leading role in the latest analysis. Researchers have spun off roads and driven into ditches; they've hit deer and, reportedly, grazed a bald eagle.

If they make it to Soudan in one piece, they stay for a week or two at a time in a gray, four-bedroom house five minutes from the mine. Every morning at 7:30, they squeeze into the mine shaft's rusty cage, where an operator lowers them into the lab. They don't return to the surface until 5:30 in the evening. The lab is within a four-story high cavern, and with walls, bright lights, a steady temperature of 70 degrees, and even a ping-pong table, it isn't too much different from a lab above ground.

Because they're waiting for a single particle among billions, at a rate of just a few potential hits per year with the current detectors, the team must minimize as much background noise as possible. With almost 700 meters of dirt and rock above, CDMS is protected from cosmic rays, which are not actually rays but primarily particles—protons, helium nuclei, electrons, and muons. A cosmic-ray strike anywhere in the instrument would trigger a cascade of other particles such as neutrons, which could mimic a WIMP's signal if they hit the detector.

Developed by groups at Stanford and UC Berkeley in the 1990s, led by Blas Cabrera and Bernard Sadoulet, respectively, the detectors are composed of 230-gram germanium crystals. The experiment's detection method requires a semiconductor, and of them all, germanium has the heaviest nuclei, maximizing the possibility of a collision. The crystals are arranged into five stacks of six detectors each, and then the stacks are encased in copper, chosen for its low radioactivity. The stacks are kept at a frigid 40 millikelvins—that's 0.04 degrees above absolute zero, almost a hundred times colder than outer space—and surrounded by layers of polyethylene, for more neutron protection, and lead, to shield against gamma rays. Because the lead will inevitably have some radioactive isotopes, spouting the occasional particle, the detectors are surrounded by ancient lead taken from an old sunken ship. Two centuries old, this lead has already decayed away most of its radioactivity. Surrounding the apparatus is something called a plastic scintillator anticoincidence detector, a shield built by UC Santa Barbara that lights up when muons pass through, notifying researchers of events to ignore. Finally, the whole thing sits in a clean room.

When a WIMP smacks into the disk-shaped detector, it rattles a germanium nucleus, sending phonons—quantized packets of vibrational energy—across the crystal. On the detector's surface are thin aluminum films, which act as phonon-collecting antennas. When aluminum is as cold as these detectors are, its electrons—which freely roam in the metal—arrange themselves in so-called Cooper pairs. These pairs form when a free electron tugs at the positively charged lattice of the aluminum nuclei, creating a tiny ripple in the crystal structure. The displacement caused by that ripple then pulls at another free electron, coupling it with the first. The partnership is weak, so the phonons from the crashing WIMP easily split the pairs, transferring their energy as heat to the newly single electrons. The electrons then find their way toward tungsten thermometers.

The tungsten is a transition-edge superconductor, meaning that any bump in temperature would nudge it out of its superconducting state. The electrons—now

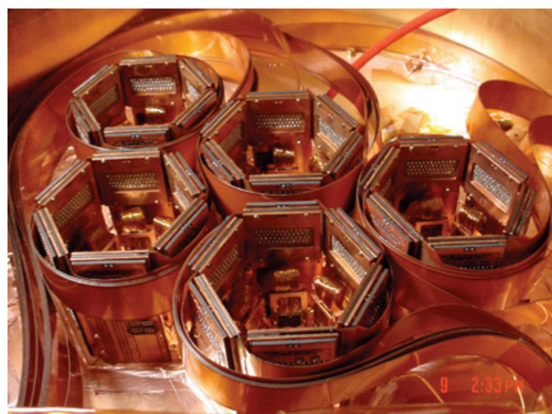

Top right: A top view of the five stacks that make up the CDMS experiment. Copper provides a first layer of protection against background particles. Bottom right: This model of the apparatus provides a multisensory demonstration of CDMS in action. It lights up and chimes when it detects particles.

Far left: The Soudan mineshaft in winter.
Middle: The park's mine tour includes a cart ride.
Left: Inside the CDMS lab, where Rupak Mahapatra (left), now on the faculty at Texas A&M University, and Xinjie Qiu (right), now a Stanford postdoc at Fermilab, take a break from dark-matter hunting with a game of ping pong.
Below: Dan Bauer, from Fermilab and project manager for CDMS, removes one of the towers to make room for the first of the next generation of towers, called SuperCDMS. With him are Jim Beaty of the University of Minnesota (center) and Steve Leman of MIT.
Below right: Walt Ogburn is in the "cryopad," the area that houses the liquid-helium and nitrogen dewars and pumps. The computer controls the circulation of the liquids that keep the detectors frigid.

warm from the phonons—push the tungsten toward being a normal conductor. By measuring the sudden changes in electrical current, the researchers can determine the energy of the WIMP-generated phonons—and therefore the energy of the collision.

To be sure, one of the biggest challenges for the CDMS team is to minimize and understand background particles. Despite the mine's depth, the shielding, and the effort to keep materials clean and pure, stray particles still slip through. Electrons from the beta decay of radioactive elements that found their way onto the detector surfaces, and high-energy photons called gamma rays from impurities in the apparatus, both knock electrons loose in the germanium. To distinguish these events from the WIMP-nucleus collisions investigators are looking for, the detector collects and counts the free charge created by the loose electrons. An electric field in the crystal corrals the free charges onto the detector's rear surface (opposite the tungsten and aluminum), where electrodes measure the ionization energy. By calculating the ratio of ionization to total collision energy, the researchers can identify which events were electron collisions and discard them.

The team employs several more tricks to distinguish WIMPs from background particles. They only look at the signals from the most probable energy range for WIMP collisions. They discount anything that happens at the edges of the detector, where signals can get fuzzy. They only consider single-collision events—since a WIMP interacts so weakly, it would only hit one atom in one detector. If the researchers see some particle plowing into multiple nuclei down the whole stack, they know it can't be a WIMP. The team also throws away any collisions that occur at the crystal's surface, as those tend to be caused by stray particles from the apparatus itself—no matter how pure your materials, some radioactive isotopes are bound to sneak in.

To prevent bias on their part—or maybe wishful thinking—the team uses the above requirements to decide on what constitutes a signal before they look at the data. Calibrating with particles generated from radioactive sources, the team determines a set of criteria that allows a reasonable number of WIMPs, but rejects virtually all background particles. Finally, they "unblind" their data by looking at it and seeing whether any fit the criteria, revealing whether any of the elusive particles have been caught.

FINDING DARK MATTER

WIMP fever was running high on December 17, when physicists packed into auditoriums in California and Illinois to hear what, if anything, CDMS had discovered. *The Economist* had written on that day, "If the rumors are true, a solution to one of the great problems of physics may now be within reach." JoAnne Hewett, a particle physicist at SLAC, even liveblogged the event on Cosmic Variance. "The excitement in the air is palpable," she wrote. "Not much work is being done—everyone is pretty much talking in the hallways, trying to pass the time until 2:00."

Finally, the results were announced—two events had been found! The team had actually unblinded the data set a month before, while Golwala was at the Caltech Submillimeter Observatory on Mauna Kea in Hawaii, where he was doing astronomical observations with Bolocam for another project. "I was in my dorm room recovering from altitude sickness that morning," he recalls. "I was floating in and out of the teleconference as the unblinding was happening. I was lying on the bed and heard something about two events. I thought, 'Uh oh. What did we do wrong?'"

After checking the data and instruments, the team concluded that nothing had gone wrong. These two collisions were real. But,

The barn was born during the Manhattan Project, when American physicists needed a shorter way of saying 10^{-24} square centimeters, which is about the size of a uranium nucleus. In the July 1972 issue of *Physics Today*, Marshall Holloway and Charles Baker recount how they coined the term in 1942, after mulling over other candidates like "Oppenheimer" or "Bethe," in honor of the physicists who made the atomic bomb possible. But "Oppenheimer" was too long and "Bethe" was too similar to the Greek letter. In no small part because of where they both were from—Purdue University in Indiana, surrounded by farmland—they decided on the "barn," since a uranium nucleus was as "big as a barn." For security reasons, physicists also wanted to avoid discussing overtly technical terms on the phone. The U.S. government ended up classifying the term anyway, and didn't declassify it until 1948. "Zepto" is the prefix for 10^{-21}, and so a zeptobarn is 10^{-45} square centimeters.

before booking flights to Stockholm, they calculated that there was a 23 percent chance these signals were caused by background—likely electron recoils that had snuck past their set of criteria. As Golwala points out, "No one claims discovery with that high of a chance." The team couldn't say they had discovered dark matter, but they couldn't rule it out, either.

So CDMS hasn't revolutionized our understanding of the universe—yet. As for all the hype? "In a couple of months, no one will remember this," Golwala says. Still, their results—published in the March 26 issue of *Science*—are noteworthy, placing the most stringent constraints yet again on what WIMPs could be. "It's an exciting time in the bigger sense, because we've been producing results from this experiment for about five years," he says. "We've been the premier experiment in this field."

These data sets are marking the end of the current chapter in dark-matter searches. But new experiments are already under way, including a rival project called XENON100, which uses liquid xenon in lieu of germanium crystals. Situated in Italy's Gran Sasso National Laboratory, the experiment is likely to surpass CDMS's results with new data soon. In the next couple of years, half a dozen more projects will begin—and they'll be many times more sensitive than CDMS. One of those is the next generation of CDMS, called SuperCDMS.

The bigger a detector's crystal, the better the chance a WIMP will hit it. Boasting five "supertowers" with a total of 15 kilograms of germanium, SuperCDMS will improve that chance about fourfold. The team will install these new detectors in Soudan as early as this year. "The experiment could legitimately detect something in the next few years," Filippini remarks.

A couple more years down the line, the team will build even larger detectors—70 units for a total of 105 kilograms. With a twentyfold boost in sensitivity, the team will need to move SuperCDMS to SNOLAB, a mine near Sudbury, Ontario, just north of Lake Huron. At around two kilometers in depth, it's the world's deepest underground laboratory.

If that weren't enough, researchers hope to bury 1.5 tons of germanium in DUSEL, the Deep Underground Science and Engineering Laboratory, which is slated to be built about 2.5 kilometers underground in South Dakota. The experiment won't happen until 2017 at the earliest, but these are the kinds of projects that will increase WIMP sensitivities by a hundred times over the next decade, making physicists optimistic about the future. "It's very possible that in the next five years we might be talking about WIMP astronomy, rather than just trying to detect something," says Filippini.

> "It's very possible that in the next five years we might be talking about WIMP astronomy, rather than just trying to detect something."

Now that the latest data run is complete, the CDMS team is turning toward making better detectors. For example, Moore is working on a new sensor design with improved resolution. Based on technology developed by Jonas Zmuidzinas (BS '81), professor of physics, these so-called microwave kinetic inductance detectors would allow researchers to better pinpoint a WIMP's crash site. Ahmed, whose dissertation will include a lot of the latest results, is now designing a device to measure minuscule amounts of radioactivity. As detectors become more sensitive, materials need to be even cleaner, with as little radioactivity as possible. Current instruments, however, can't detect such low levels.

WIMP-search experiments are approaching what physicists call the zeptobarn scale (see box). The zeptobarn, which is 10^{-45} square centimeters, is a unit of cross section, the measure of how frequently a particle interacts. The larger the cross section, the more likely those particles will collide. The latest CDMS results rule out WIMPs with cross sections bigger than about 10^{-44} square centimeters, or 10 zeptobarns. If

WIMPs have cross sections bigger than this limit, then CDMS should've already detected them.

It turns out that the most compelling theories of supersymmetry or extra dimensions predict WIMPs to have cross sections of around a zeptobarn. Because the next generation of WIMP detectors will be sensitive to these kinds of particles, they could be discovered in the coming decade, Golwala says. But below a zeptobarn, things get weird. If you don't see anything by a zeptobarn, "you're excluding the most reasonable models," he says. "But no one said nature had to be reasonable." At cross sections of hundredths or thousandths of a zeptobarn, the likelihood of dark matter being WIMPs drops, and with limited funding, it might not be worthwhile to continue the search.

Regardless of when scientists discover dark-matter particles, it likely won't happen with a single, big announcement in one-inch headlines—despite what rumors might lead you to believe. The search is a slow, systematic sweep of the possible identities of dark matter, and any finding will have to withstand close scrutiny and be confirmed by multiple experiments.

Direct-detection experiments like CDMS are just one of three roads to determining whether dark matter is made of WIMPs. Physicists are trying to detect WIMPs indirectly with satellites like Fermi, which aims to measure the gamma rays that are shot out from WIMP-WIMP annihilations. They also hope to make WIMPs from scratch at the Large Hadron Collider in Geneva, which will be able to recreate the conditions of the universe moments after the Big Bang. These experiments represent a confluence of esoteric theory and tangible hardware that may soon solve one of the great mysteries of nature. "It's amazing that we can even ask the questions, what is the universe made of, and why it's made of that stuff," Golwala says. "And it's even more amazing that we can attempt to answer them." e&s

EUREKA! . . . OR NOT

There have been plenty of tantalizing "discoveries" of dark matter. One experiment in Italy, called DAMA, has claimed to see dark matter not once, but many times. DAMA, like CDMS, tries to measure WIMP collisions. Instead of germanium, DAMA has sodium iodide crystals that glow when struck by particles, and instead of picking out individual WIMPs, it tries to measure the periodic ups and downs of collisions as Earth sails through the dark matter in our solar system. This annual modulation, as it's called, could be a signature for WIMPs. The group said they detected this signature first in 1997 and that they confirmed the signal multiple times, in 2000, 2004, and 2008.

The problem, though, is that no one else has been able to confirm the results. You have to be sure the signal isn't coming from another, non-WIMP source, such as an exploding star or just some artifact of your technique. "Over the last decade, DAMA has managed to slowly knock down all the objections I have had to their result, in terms of possible non-dark-matter explanations," Golwala says. "But that doesn't mean that all non-dark-matter options have been exhausted. And they have never been particularly forthcoming with their data, in spite of the extraordinary claim to have detected dark matter. Again, there is no way to conclusively prove it, but I tend to think it's not dark matter."

Another way to detect WIMPs is to see them annihilate one another in space. In the moments following the Big Bang, pairs of WIMPs—if WIMPs actually exist—would have been smashing into each other, igniting bursts of energy and showers of particles. Even though they wouldn't be crashing as frequently now as they did then, they would still manage to find one another once in a while. In 2008, a satellite called PAMELA (Payload for Antimatter Matter Exploration and Light-nuclei Astrophysics) and a balloon over Antarctica, ATIC (Advanced Thin Ionization Calorimeter), found complementary data that could be explained by WIMP annihilations. ATIC detected an excess of high-energy electrons, while PAMELA found extra high-energy positrons, or antielectrons. Last year, Fermi, NASA's gamma-ray space telescope, also detected a subtle boost in electrons.

But these extra particles could also have come from other sources, such as a type of dense rotating star called a pulsar. Golwala remains skeptical. "It's not possible to say this conclusively, but there have been too many of these signals that appear and then can be explained astrophysically to put much stock in any of them." All the different data don't fit together in a coherent picture, adds Moore. "There's no real smoking gun for any of those detections yet." e&s

Sunil Golwala received his BA in physics from the University of Chicago in 1993. He received his MA in 1995 and PhD in 2000 from UC Berkeley, where he worked on the previous generation of the CDMS experiment with Bernard Sadoulet. He came to Caltech in 2000 as a Millikan Postdoctoral Scholar, became an Assistant Professor of Physics in 2003, and Associate Professor of Physics in 2010.

His CDMS work is supported by the Department of Energy (DOE). The current CDMS experiment was built by the DOE and the National Science Foundation.

Cassini’s Ringside Seat

Saturn’s rings and moons are a model for how planetary systems may be forming around nearby stars. The Cassini spacecraft has been taking a close look at the rings of late, and it turns out that gravity and granules with momentum can do some amazing things.

Of all the planets in our solar system, only one’s got a whole lot of bling. Yet, for most of human history, we didn’t know about Saturn’s luminous rings, even though the gigantic, gaseous planet is readily visible in the night sky. When Galileo became the first person to peer at Saturn through a telescope in 1610, he sketched a companion moon on each side. A couple of years later, he became utterly perplexed when the moons vanished. When the mysterious objects returned, he saw them as elliptical arms resembling handles. In 1655, Dutch astronomer Christiaan Huygens, using a better telescope, figured out that the arms were actually a flat ring and that, like the edge of a sheet of paper held horizontally at eye level, it disappeared from view when Saturn’s tilt presented it to us edge-on. (The rings truly are paper-thin—a mere 10 to 20 meters thick on average, and yet so broad that they would fit neatly between Earth and

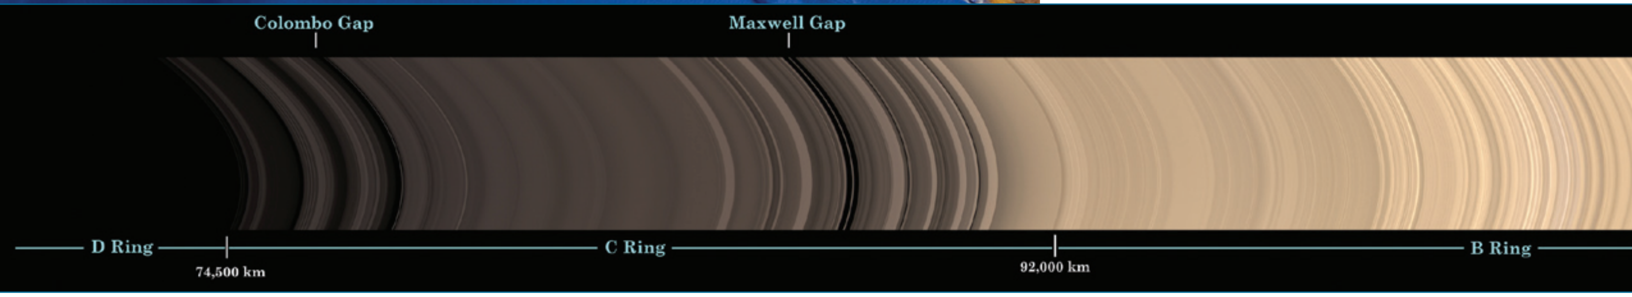

Opposite: Linda Spilker, Cassini's project scientist, with one of the Deep Space Network's 34-meter dishes at Goldstone, California.
Right: The Cassini orbiter.
Below: The anatomy of Saturn's rings. The planet itself is out of view to the left.

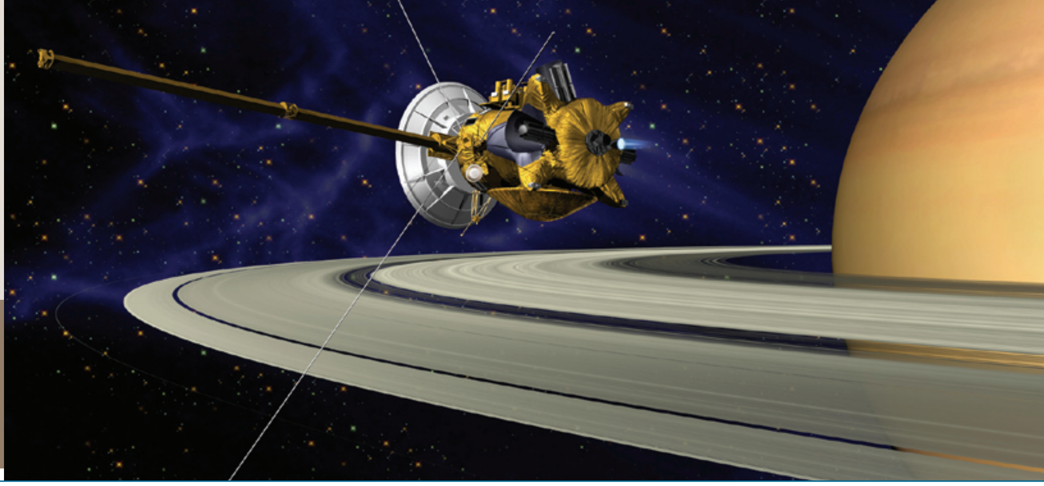

By Linda Doran

its moon.) Saturn's rotational axis is tilted by about 27 degrees, and the rings, which girdle Saturn's equator, are tilted by the same amount. Seasons change slowly on Saturn, with its nearly 30-year orbit; but once every 14 to 15 years, when the sun straddles Saturn's equator at the solar equinox, the rings vanish from view—they're pointed directly at the sun and more or less at us.

Seen through the 60-inch telescope at Caltech's Palomar Observatory, Saturn's rings beckon like diamond-paved, circular highways. (Every now and then, the Friends of Palomar Observatory get a tour that, weather permitting, includes an opportunity to turn that telescope on the heavens.) They nearly fill the telescope's field of view, and seem as close as the moon does in a pair of good binoculars. In reality, they're more than a billion kilometers away. If you were to count those kilometers at one per second, you'd still be counting three decades from now—an entire Saturnian year later. Put another way, Saturn is about 10 times farther away from the sun than Earth. Understanding something so remote is a challenge not only for parents with curious children—at Palomar and elsewhere—but for scientists as well. To this day, we're not even sure how old the rings are or exactly what they're made of.

Our first clue came in 1675, when Giovanni Domenico Cassini, the founding director of the Paris Observatory under Louis XIV, observed not one bright ring but two, separated by a dark gap. The discovery of what's now called the Cassini Division, 4,800 kilometers wide, demonstrated that Saturn's rings were not a single, solid object. Quite perceptively, Cassini theorized that they were swarms of tiny moonlets too small to be seen from afar.

Three centuries later, spacefaring vehicles carried telescopic cameras to Saturn itself. Our first emissary, NASA's Pioneer 11, revealed in 1979 that the Cassini Division was not actually empty, but was merely a zone with a very low density of ring material. JPL's twin Voyager spacecraft, flying past Saturn in 1980 and 1981, showed that the rings are astonishingly complex. Astronomers and stargazers alike marveled at images of gaps swept clean by tiny moons, mysterious dark spokes whirling within the B ring as if it were a wagon wheel, and the "braided" multiple strands of the F ring. The Voyagers showed that Saturn's rings were really thousands of discrete ringlets, and confirmed that they are very nearly pure—99 percent—water ice. (We'll come back to that other 1 percent later.) The Voyagers also determined that the ring particles range from flecks the size of dust, about a millionth of a meter across, to chunks of ice as big as two-story houses. Most of them, however, run from pebbles to snowballs.

On June 30, 2004, a schoolbus-sized spacecraft named Cassini, a joint U.S.-European mission built and operated by the Jet Propulsion Laboratory, which Caltech manages for NASA, entered orbit around Saturn. Six months later, on December 25, Cassini released the European Space Agency's Huygens probe, which parachuted to the surface of smog-shrouded Titan. Discovered by Christiaan Huygens in 1655 and bigger than the planet Mercury, Titan is the second-largest moon in the solar system and has an Earth-like surface with highlands, lowlands, stream channels, and liquid methane–filled lakes. Huygens sent pictures and data back to Cassini during the entire journey down, and then continued to transmit information from the surface for some 70 minutes, until the mother ship flew out of range. Meanwhile, Cassini had settled into its four-year primary mission—a grand tour of a planet that many scientists regard as a miniature solar system in its own right. A team of navigators guides Cassini on the multiple loops of its itinerary, firing rocket thrusters to fine-tune its orbit and using the gravitational pull of Titan like a slingshot to adjust the spacecraft's path. Cassini's travels have taken it on equatorial orbits in the plane of Saturn's rings and moons, as well as on steeply inclined orbits above and below the plane in order to make observations from as many angles as possible. And Cassini will be going strong for years to come—in January 2010, NASA authorized a seven-year mission extension to follow Saturn through a complete change of seasons.

PLUNGING INTO THE DARK

Which brings us back to the solar equinox. In August 2009, for the first time in history, astronomers got a close-up look at Saturn's rings shrouded in darkness, just as they were when Galileo couldn't see

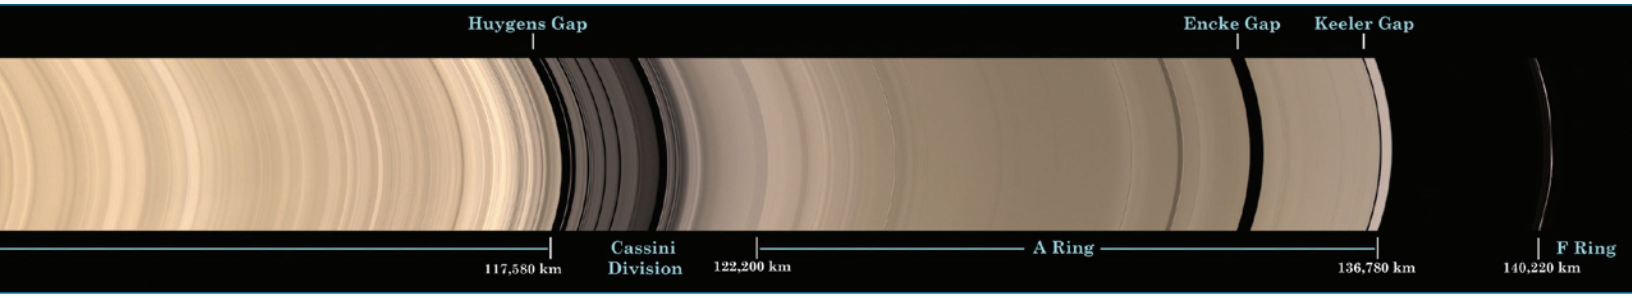

The gravitational pull of a tiny moon named Daphnis, only eight kilometers in diameter, sweeps stray material out of the 42-kilometer-wide Keeler Gap near the outer edge of the A ring. Cassini discovered that Daphnis's gravity also kicks up bright, white rooster tails in the ring material.

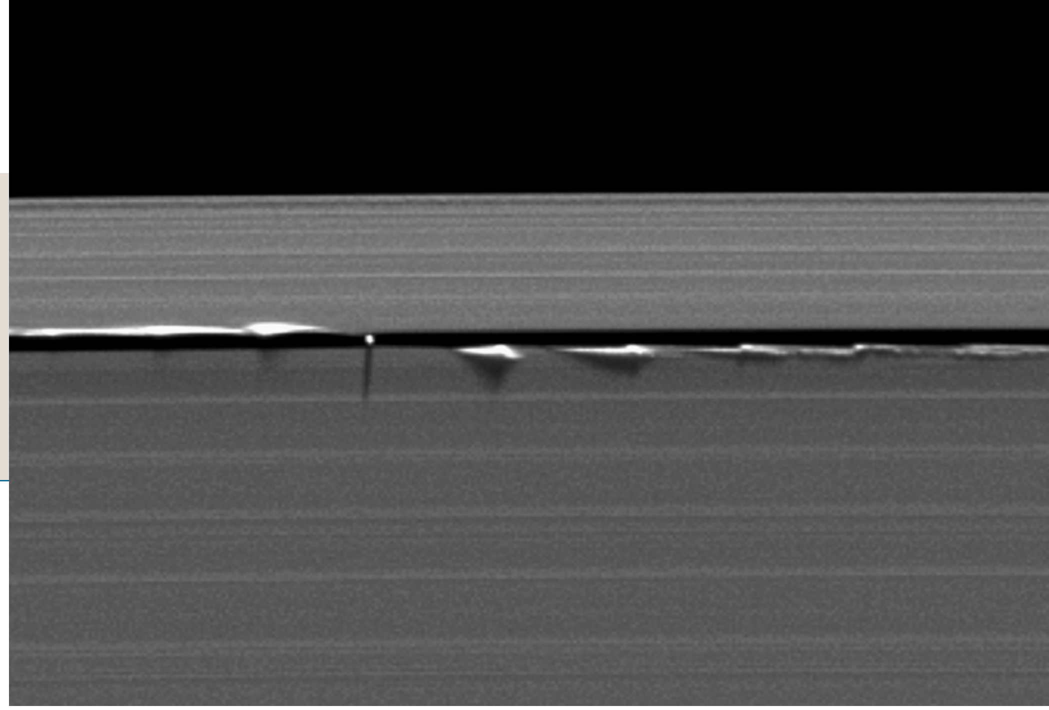

them. One of those astronomers was JPL's Linda Spilker, Cassini's project scientist. "At equinox, the sun is essentially edge-on to Saturn's rings," she says. "This is the time when the ring temperature is at its very coldest—down to 43 kelvins (–230 degrees C). In a sense, it's like experimentally turning off the sun. The rings are heated only by Saturn's thermal radiation and by sunlight reflected from Saturn." When Cassini arrived at Saturn, the rings averaged 90 to 110 K (–183 degrees C to –163 degrees C). During equinox, a 60-degree temperature drop happened overnight, as it were.

Spilker belongs to that generation of space explorers who spent July 20, 1969, in wide-eyed wonder in front of their television sets watching Neil Armstrong set foot on the moon. While an undergrad majoring in physics at Cal State Fullerton, she studied other space rocks at Caltech—helping her physics professor, Dorothy "Dotty" Woolum, who had been a postdoc at Caltech in the '70s, and Don Burnett, Caltech professor of nuclear geophysics, analyze the distribution of bismuth and lead in meteorites. Later, she worked on the Voyager mission, and she likes to tell her daughters that their births were literally based on the alignment of the planets. The girls were born between Voyager 2's flybys of Saturn in 1981 and Uranus in 1986—in other words, during the "cruise phase" when there wasn't so much going on.

Spilker's fascination with destinations beyond Earth is a constant in her life. "If you could scoop up a particle out of Saturn's rings," she asks, "and hold it in your hand, what would it look like? Would you see this fluffy snowball? Would it have an icy core? If you could take your knapsack and collect particles as you collect shells on a beach, would you find different kinds of particles in each of the rings? I like being an explorer."

Spilker belongs to that generation of space explorers who spent July 20, 1969, in wide-eyed wonder in front of their television sets watching Neil Armstrong set foot on the moon.

She has a partial answer to those questions as a result of data collected during the equinox. In addition to her management role coordinating Cassini's many investigations, Spilker is coinvestigator on an instrument called the Composite Infrared Spectrometer, or CIRS. Humans sense infrared light as heat. Like a snake's tongue, CIRS seeks out sources of heat, experiencing the world at wavelengths longer than those the human eye can see. As it distinguishes variations in infrared light, CIRS can map the temperature and composition of Saturn's rings because different materials emit and absorb heat in characteristic ways.

"I think that the outsides of Saturn's ring particles would certainly be very fluffy and porous," Spilker says. "That's based on the fact that they heat up and cool down very quickly, within about half an hour. They have what we call a very low thermal inertia—in fact, it's four orders of magnitude lower than that of a solid block of ice." Rapid heating and cooling implies a lot of surface area in order to exchange so much thermal energy with the surroundings so fast. That's why crushed ice in a glass of soda melts faster than cubed ice.

For about four days, the rings were cloaked in shadow. Indeed, sunlight on Saturn's rings during the equinox is so dim that the rings are nearly invisible in the raw camera images, even those taken from overhead. The only reason we can see them is that scientists led by Carolyn Porco (MS '79, PhD '83) digitally enhanced the shots from Cassini's wide-angle camera. (Porco, too, worked on the Voyager mission, and is now also working on NASA's New Horizons mission en route to Pluto.) To bring the rings to life, she and her colleagues at the Cassini Imaging Central Laboratory for Operations (CICLOPS) at the Space Science Institute in Boulder, Colorado, increased the brightness of the dark half of the rings by a factor of three relative to the half illuminated by Saturn-shine, and bumped up the brightness of the entire ring system by a factor of 20.

The photos revealed walls of icy rubble as tall as mountain peaks rising out of the ring plane. "It's like standing outside right before the sun sets," notes Spilker. "Your shadow gets very long. Anything that's a little bit bigger, or sticks up, casts shadows on the part of the ring that's behind it. We see what look like towering mountains, in some cases as high as four kilometers, created by the particles that Saturn's moon Daphnis is pulling out of the ring plane." This happens because Daphnis's orbit is slightly tilted with respect to the rings. When Daphnis crosses the ring plane, it drags some of the particles out of the plane. Like the rooster tail behind a speedboat at full throttle, this disturbance quickly subsides once Daphnis passes by. If you were a pilot flying above the great plains that are Saturn's rings, it would be a good idea to keep plenty of altitude over the landscape below, lest you crash into an icy curtain as tall as the Rockies.

As you kept an eye out for ice peaks, you might also notice regularly spaced basins and ranges. Close analysis of the Voyager pictures had revealed spiral density waves in the rings. These waves occur wherever one of Saturn's moons is in resonance with a region in one of Saturn's rings. Orbital resonance brings two bodies back to the same region of space over and over again—for example, if the ring particles travel around Saturn exactly twice for every orbit of the moon, they have a 2-to-1 resonance. The repeated gravitational tug at that location causes the ring particles to crowd together into a coherent spiral structure one or two kilometers wide. The resulting density wave sweeps through the ring material. As the crest overtakes the ring particles in their orbits, they get pushed together by the wave; once the wave passes, the trough pulls the particles back into their original positions relative to their fellows. These density waves are analogous to the arms of a spiral galaxy such as our own Milky Way, says Porco. "They're much more tightly wound than the spiral structures that you see in galaxies, but they are in fact the same creatures. The physics behind them is the same."

Spiral density waves are compressional, existing only in the two dimensions of the ring plane. But if the orbit of the resonant moon is slightly askew, yet another set of waves forms—this time, in three dimensions. The orbit of Saturn's moon Mimas also takes it above and below the ring plane, pulling ring particles out of the plane along

Planetary scientist Carolyn Porco did her PhD thesis on Voyager's observations of Saturn's rings and spokes. She now leads Cassini's imaging team, which is headquartered at the Space Science Institute (SSI) in Boulder, Colorado. In 2009, she received the most prestigious award in scientific photography, the Lennart Nilsson Award, "for combining the finest techniques of planetary exploration and scientific research with aesthetic finesse and educational talent."

"From the very beginning of the Cassini mission," says Porco, "I took it as my personal goal that we on the imaging team would be the planetary equivalent of nature photographers. We would try to capture, whenever possible, images and movies whose primary purpose went beyond science and conveyed the sheer magnificence and otherworldly beauty to be found around Saturn."

The vast library of aesthetically striking images taken during equinox as well as throughout the mission is available for viewing at JPL's Cassini website and at the website of SSI's Cassini Imaging Central Laboratory for Operations (CICLOPS), where the images are processed. The images paint a portrait of a planetary system that, in Porco's words, "is so alien in comparison to Earth that we might as well have visited a planet around another star in another quadrant of the galaxy. When this mission is over, we will leave behind a stunning visual legacy, and a body of work that will guide future explorers, both robotic and human, in their excursions around the Saturn system."

A much sought-after speaker and Carl Sagan protégée who served as a consultant on the movie *Contact*, she also consulted with Industrial Light and Magic's visual-effects supervisor Roger Guyett on the recent *Star Trek* reboot. This picture, taken at ILM, shows her with some familiar faces from yet another blockbuster franchise. (Photo courtesy of Carolyn Porco and ILM.) E&S

with it. Orbital resonance then keeps these particles' orbits tilted relative to the plane, resulting in a towering spiral wave called a bending wave that wraps all the way around the ring. The density waves and the bending waves propagate in opposite directions, says Spilker, adding a further level of complexity to the rings' structure.

And now Cassini's equinox pictures show yet another kind of three-dimensional wave, heretofore undiscovered. These waves only rise about 100 meters above the ring plane, yet during the equinoctial photo session they too remained illuminated after the rings were plunged into darkness. "From high above, the rippled surface of Saturn's D ring looks like a corrugated roof," says Spilker. "The ripple extends for more than 17,000 kilometers across the ring system. When the Voyagers flew by, it wasn't there. Something caused this wonderful rippling to get started, and then it expanded outward all the way into the C ring. Based on computer models of the ripple's expansion, we can run the clock backward, and we think that it began some time in the early to mid-1980s." Scientists remain baffled about its origin, but one possible scenario has a meteoroid slamming into the rings—the pebble in the pond, as it were. Like the waves the Voyagers discovered, this ripple is also a tightly wound spiral. In fact, it has become more tightly wound over time. In Hubble Space Telescope images taken before Cassini's discovery and since reexamined, the ripples were farther apart.

Cassini also got a close-up look at the mysterious "spokes" in Saturn's rings first observed by the Voyagers. These spokes may be the result of meteoroids hitting the ring and generating charged particles, says Porco. The particles levitate, in the same manner that a comb rubbed vigorously on a wool sweater will cause your hair to rise if held near your head, until eventually they lose their charges and settle back into the ring plane. The low-angled sunlight at equinox caught the high-flying spokes, setting them off against the shadowed rings below.

BLINKING IN THE SUNLIGHT

If, on a bright sunny day (you certainly wouldn't want to be using visual flight rules during equinox), you were to buzz down between the A and B rings through the Cassini Division, you would find that the rings' unlit side is not completely dark. Instead of being continuous, each of Saturn's ringlets is like a dashed line of clouds. "As we fly underneath," says University of Colorado at Boulder planetary scientist Larry Esposito, "we see flashes of sunlight come through the gaps. We originally thought we would see a uniform cloud of particles. Instead we find that the particles are clumped together with empty spaces in between." Esposito is principal investigator for the Ultraviolet Imaging Spectrograph, or UVIS, which sees light at wavelengths shorter than we can perceive—ultraviolet rays are what burn your skin if you forget the SPF 30. Esposito's team measured ultraviolet emissions from a distant star, Alpha Arae, which waxed and waned as the rings passed in front of it.

Cassini confirmed the presence of the taffylike clumps of material by sending radio signals through the rings to Earth. These signals—with a strength of less than a billionth of a watt by the time they arrived here!—were picked up by the ultrasensitive receivers of JPL's Deep Space Network, which has stations in California, Spain, and Australia in order to ensure that one set of ears is above the horizon at all times. Besides communicating with NASA's robotic explorers (and spacecraft of other nations as well), the 70- and 34-meter dishes are also used for radio-science experiments. In this case, the ring particles absorbed and scattered the radio waves in a manner that revealed their size and distribution.

The clumps, called self-gravity wakes, are a special case of the spiral density waves mentioned earlier, which Cassini has now shown to exist all over Saturn's rings. "We first saw self-gravity wakes in the A ring, which is less dense, so they are farther apart and easier to see. They just leaped out at us," says Spilker. "They are harder to tease out in

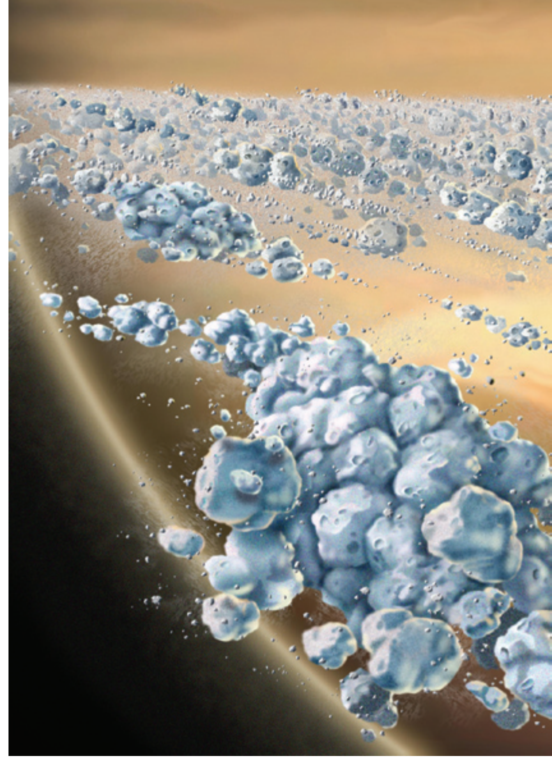

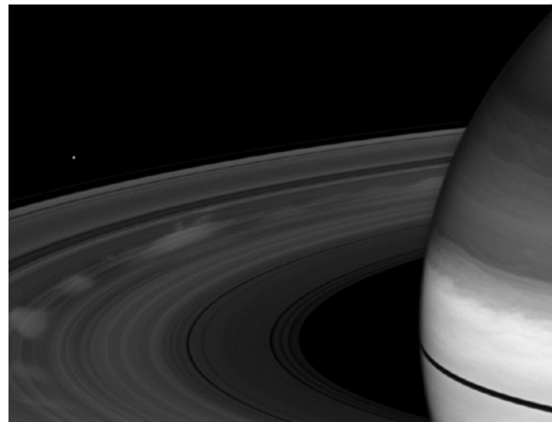

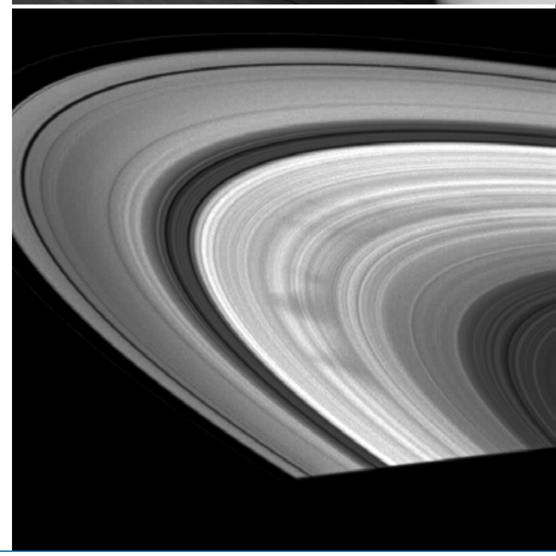

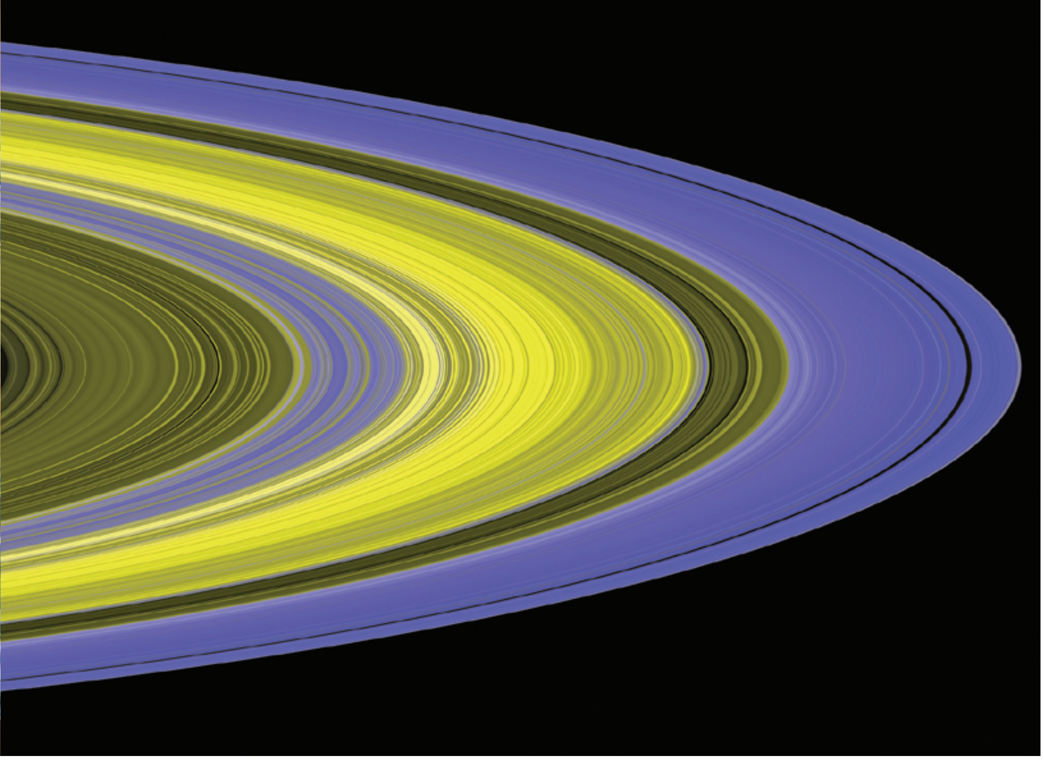

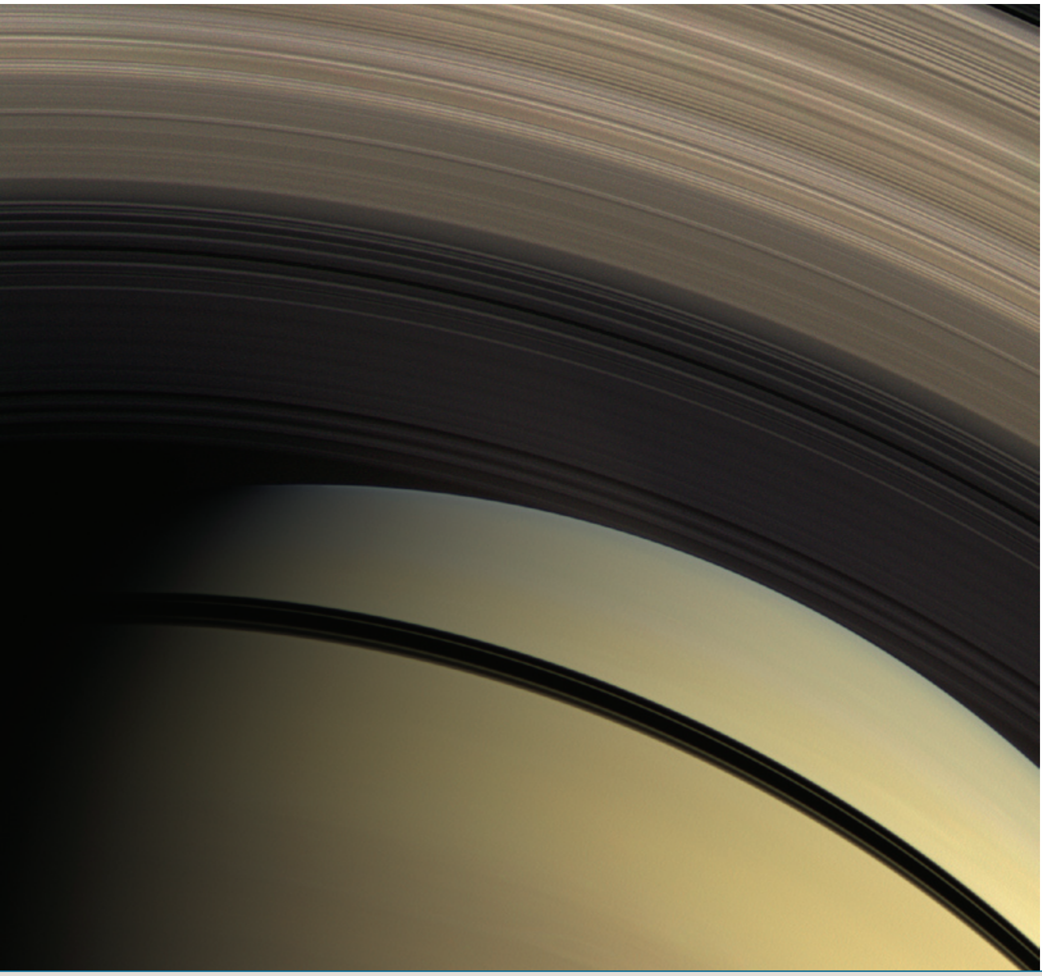

the B ring, because it is so optically thick."

All the particles in all the rings are constantly colliding, but the dense B ring affords the most chances for the particles to adhere to one another afterward, held together by gravity and their own stickiness. The resulting flat sheets of material grow by accretion until they are 30 to 50 meters wide, at which point Saturn's gravity pulls them apart. Under different circumstances, if they were farther from the planet, the wakes might have been the seeds of moons.

And what of those ring particles? What are they made of? Though we know that water ice is 99 percent of the ring material, the other 1 percent remains a puzzle. Cassini has returned spectacular images showing that the rings when bathed in sunlight appear golden, golden-brown, or even slightly pinkish. If the rings were made only of water ice, they would be frosty white or bluish-white, like glacial ice on Earth. Something else is creating the subtle hues seen by Cassini, which has cameras with much greater color sensitivity than those on earlier spacecraft, notes Spilker.

"Originally, after Voyager 2 flew past Saturn, we thought the fact that the rings were so bright and clean must have meant that they were young. Now we're starting to think that maybe they're a mix of young and old. There may be some processes that periodically break the ring particles open and recycle the contents. You can imagine breaking open a snowball and releasing fresh material into the system, making the rings appear cleaner and brighter. We'd really like to get an idea of the rate at which the rings are accumulating additional, non-icy material and then figure out what kind of processing is going on within the rings."

Spilker thinks that soon we'll be able to tease out the subtle signals associated with ring contaminants from the very large signal that corresponds to water ice. One of the instruments working overtime on identifying the chemical makeup of the rings is Cassini's Visual and Infrared Mapping Spectrometer, or VIMS, which maps colors at

Clockwise, from top left: 1. An artist's rendition of the ring particles. They continually clump and disperse again, forming thin, curved transitory aggregates with nearly empty space in between.

2. This false-color image from UVIS shows the density and orientation of the clumps. The brightest regions are the densest—in fact, the middle of the B ring was too dense for Alpha Arae's light to penetrate. The clumps are tilted into oblique wakes in the blue regions, and are oriented along concentric circles in the yellow ones.

3. A natural-color view of the sunlit rings, as seen from below. The translucent C ring runs through the center of the frame, while the denser B ring arcs across the top.

4. As Saturn neared equinox in November 2008, the "spokes" (seen here as dark smudges) first seen by the Voyagers one Saturn year earlier returned to prominence.

5. In September 2009, a month after equinox, the spokes stand out much more clearly. Janus, a small moon 179 kilometers in diameter, can be seen at upper left.

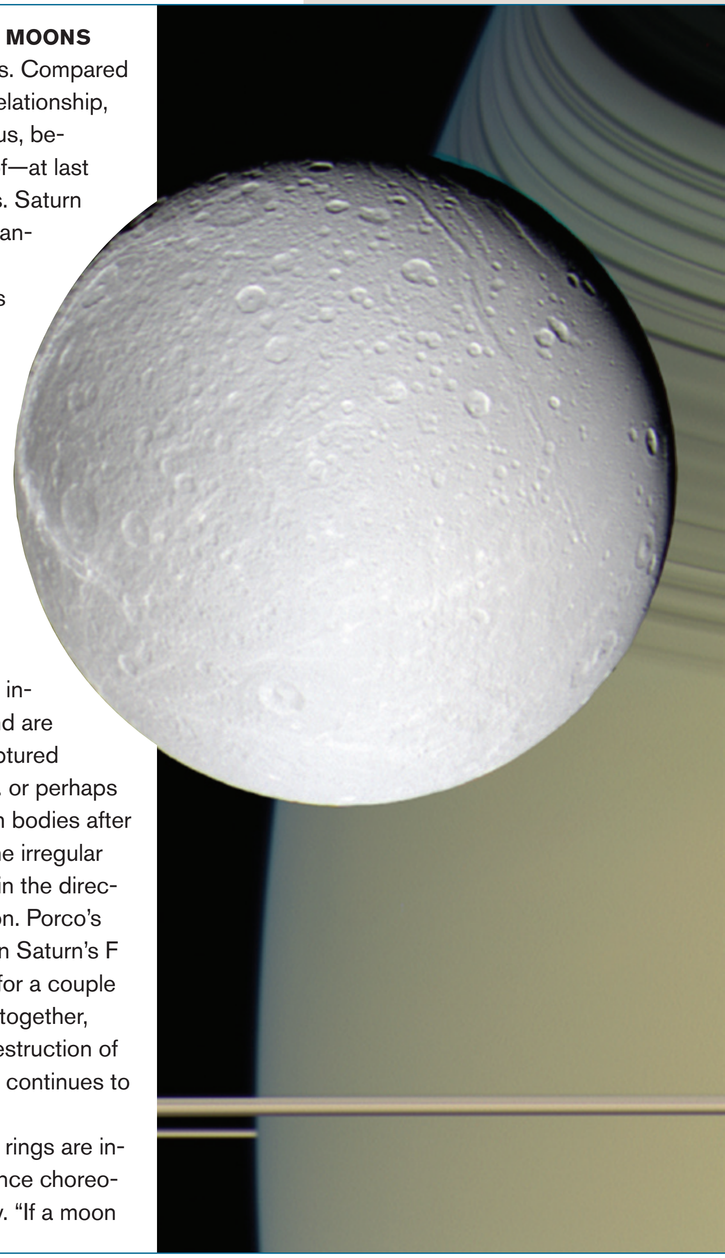

Below: A natural-color shot of icy Dione, snapped when Cassini was very nearly in the ring plane. The rings form a razor-thin horizontal stripe across the bottom of the picture, while the set of narrow, curving shadows cast on Saturn by the C ring is visible behind Dione. A portion of the B ring's shadow adorns the upper right corner.

different wavelengths than does CIRS. Like CIRS, the VIMS spectrometer breaks light into its component wavelengths. Because every molecule reflects a specific set of wavelengths and absorbs the rest, scientists can identify molecules by their spectral "fingerprints." Shortly after Cassini arrived at Saturn, VIMS data hinted at the presence of some kind of iron-bearing compound—perhaps not too surprising, as many scientists think the rings are the remnants of icy comets and iron-bearing meteoroids torn apart by Saturn's gravity.

The composition of the rings has a direct bearing on their age. Water ice is easily eroded, and ring particles are pulverized by micrometeoroids and ground up by collisions with each other. Saturn itself plays a role in breaking down the ice grains—UVIS has detected an immense cloud of oxygen atoms liberated from ring ice by bombardment from Saturn's own internal radiation. The cloud surrounds Saturn and extends for millions of kilometers beyond Saturn itself, says Esposito. The rings also have their own atmosphere of oxygen gas, produced when ultraviolet light from the sun interacts with the water ice. Other kinds of debris, such as tiny pieces of rock or minerals from micrometeorites, are harder to erode and have greater longevity. "It could very well be that different parts of the ring system have different ages," says Porco. "The massive middle B ring might be a lot older than the A ring. One might be billions of years old and the other only a few tens of millions or hundreds of millions of years. We don't know."

MOONS, MOONS, AND MORE MOONS

And then there are the moons. Compared to Earth and its monogamous relationship, Saturn is downright promiscuous, being accompanied by a retinue of—at last count—62 orbiting companions. Saturn is surrounded by real moons, wannabe moons, pieces of moons, and fleeting moons. Cassini has discovered seven moons—not to mention hints of numerous suspected but unconfirmed moons that we infer from their effects on the structure of the rings—and during roughly the same interval, other observers have found 24 more. These moons are all tiny—less than 18 kilometers wide—and irregularly shaped. They travel in eccentric, steeply inclined orbits far from Saturn, and are thought to be minor planets captured by Saturn's gravitational pull . . . or perhaps debris from the breakup of such bodies after they were captured. Some of the irregular moons are retrograde, orbiting in the direction opposite to Saturn's rotation. Porco's team even discovered objects in Saturn's F ring that could only be tracked for a couple of orbits before disappearing altogether, suggesting that creation and destruction of moonlike bodies around Saturn continues to this day.

Saturn's moons and Saturn's rings are inextricably linked in a cosmic dance choreographed by the forces of gravity. "If a moon

Compared to Earth and its monogamous relationship, Saturn is downright promiscuous, being accompanied by a retinue of—at last count—62 orbiting companions. Saturn is surrounded by real moons, wannabe moons, pieces of moons, and fleeting moons.

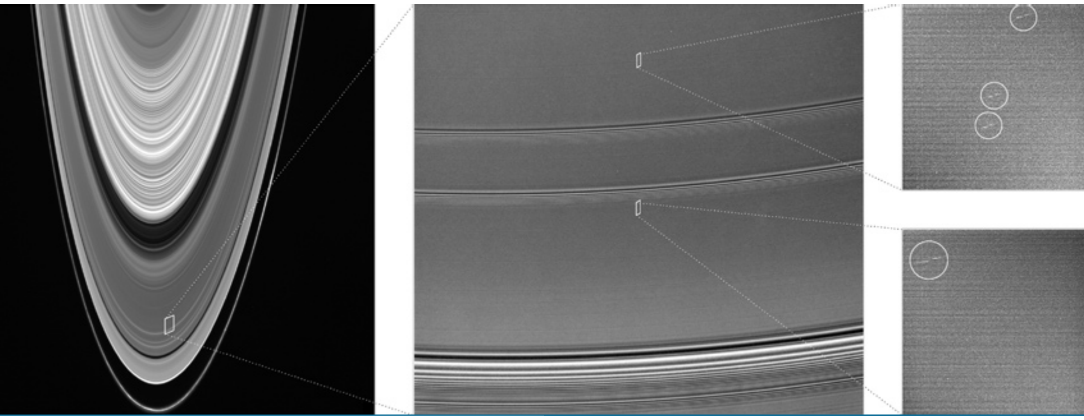

The propellers in context. The leftmost photo shows the F, A, and B rings. Zooming in on the box in the middle of the A ring (center), we see a large density wave caused by Janus and Epimetheus at the bottom of the frame and two smaller density waves in the middle. Zooming in on the featureless regions between the waves, we find four propellers (circled), each about five kilometers from tip to tip and caused by an unseen moonlet about 100 meters in diameter. The leading blade is about 300 meters closer to Saturn than the trailing one; the resolution is 52 meters per pixel.

is large enough," explains Spilker, "it can clear a gap in Saturn's rings by exerting a gravitational pull on the ring particles within a certain distance of itself. Good examples of that are two of Saturn's 'ring moons,' Pan and Daphnis. Pan orbits in the center of the Encke gap, and it keeps that gap open and pretty much free of ring particles, except for a few dusty ringlets in the gap. Daphnis, a much smaller moon, keeps the Keeler gap open. It's a much narrower gap because the moon is smaller."

In a way, Saturn's rings are a laboratory for determining how small an orbiting object can be and still be considered a moon. "Cassini found still smaller objects that would like to open a gap but don't have enough gravity to succeed," says Spilker. Instead, like speedboats, they create "wakes" that travel with them. But rather than being V-shaped, these wakes resemble two-bladed airplane propellers. One blade sticks out ahead of the moonlet, pulled forward by gravitational attraction to faster-moving ring material nearer to Saturn. The other blade stretches out behind, tugging on slower-moving material farther from the planet. The propellers can be as long as five kilometers from tip to tip. A propeller's span depends on the size of its moonlet, but even the biggest moonlets are still too small to be seen. Cassini has so far discovered more than 150 such moonlets.

Out at the edge of the main ring system lies the narrow, quirky F ring, held in place by the shepherd moons Prometheus and Pandora. Cassini has photographed the F ring in exquisite detail. We can now see that it clumps in places where aggregates of ring material or hidden moonlets lie, and kinks where the gravitational pull of Prometheus tugs on material in the ring. In fact, Prometheus literally collides with the F ring at the point in its orbit most distant from Saturn. As it moves back toward the planet and away from the ring, the moon pulls material out of the F ring, leaving dark channels behind.

In a way, Galileo was not that far off when he mistook Saturn's rings for companion moons. Sometimes Saturn's rings are a source of material for the moons, and sometimes Saturn's moons are a source of material for the rings. Cassini has confirmed that at least three of Saturn's moons are gaining girth from the rings. Pan, Daphnis, and Atlas started out as "football-shaped bodies," says Porco. "But as they sweep their paths clean, some material accretes on their surfaces, forming waistline bulges that make them look a bit like flying saucers."

Several of Saturn's moons produce rings or partial rings. The most celebrated of these is Enceladus, which spews geysers of water ice and vapor laced with organic material. (See *E&S* No. 1, 2006.) The icy particles are pulled into orbit around Saturn to form the diffuse E ring. "The geysers of Enceladus," says Porco, "are perhaps our most stupendous and significant discovery, because they very likely erupt from pockets of organic-rich liquid water. Not only does this suggest a potentially habitable environment below the south pole of the moon, but

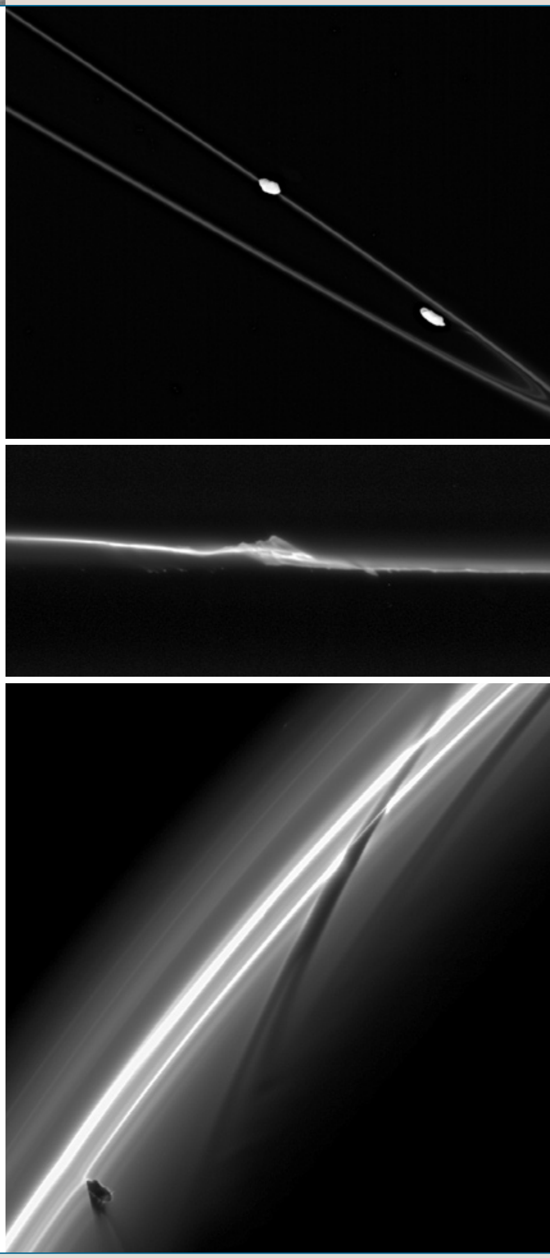

Top: Potato-shaped Pandora (left), Prometheus (right), and the F ring.
Middle: A hidden moonlet's gravity disturbs the F ring, swirling it like a wisp of smoke.
Bottom: Cassini snapped this shot half an hour after Prometheus (in the bottom left corner) burst up through the F ring—in fact, its shadow on the ring shows that it's not quite all the way out yet. It's beginning to pull a streamer of material after itself, which will eventually leave a dark channel in the ring. On its previous pass about 15 hours earlier, it tore open the channel that stretches downward across the center of the image; an even earlier channel that has started to fill back in extends into the upper right corner and out of frame.

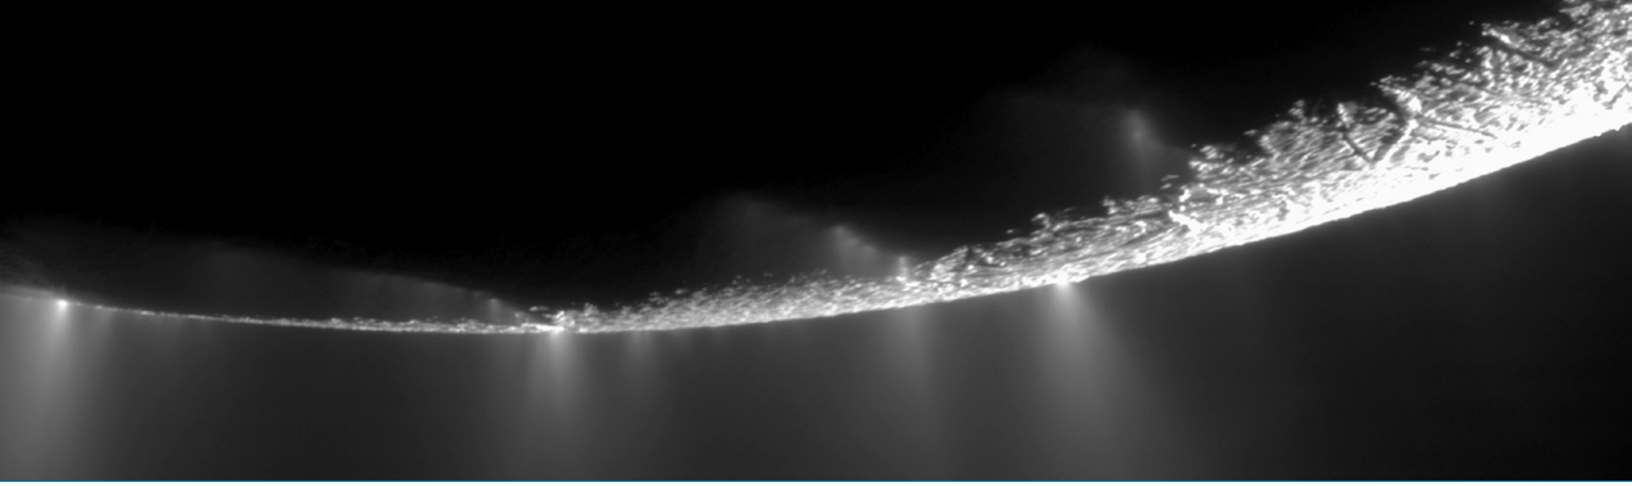

the way in which Enceladus is generating heat—and it is generating a lot of heat—is a fascinating problem in the study of planetary moons." Several other rings and ring arcs discovered by Cassini are believed to be made of dust ejected from the surfaces of various moons by meteoroid impacts. These include the Janus/Epimetheus ring, the Methone ring arc, the Anthe ring arc, and the Pallene ring.

Janus and Epimetheus, by the way, are unique in the solar system, as far as we know. The two moons essentially share the same orbit, with Janus until recently being some 50 kilometers "inside" Epimetheus. The inside moon travels slightly faster because it is slightly closer to Saturn, and since Janus is 181 kilometers across and Epimetheus is 116 kilometers wide, you'd think they'd plow into each other. Not so—when the inner moon overtakes the outer one once every four years, gravity steps in. The outer moon pulls on the inner moon, giving it extra momentum and flinging it into a higher orbit in which it, paradoxically, moves more slowly. At the same time, the inner moon tugs on the outer moon, siphoning off some of its momentum and dropping it into a lower, faster orbit. The moons trade orbits as they retreat back the way they came. The chase begins again, but now with a new pursuer. This cosmic game of tag was most recently played out this past January, putting Janus on the outside track until 2014, says Cornell's Matthew Tiscareno (BS '98), an associate on the Cassini Imaging Team. Like the grooves on a vinyl LP (remember those?), it turns out that some of the spiral density waves in Saturn's rings carry a record of the duo's past orbital swaps, as he discovered after the 2006 exchange. There is undoubtedly other information encoded in the rings as well, he says, adding that we are just starting to figure out how to extract it.

What's next for Cassini? The mission extension to 2017 will allow us to follow a complete change of seasons on Saturn, as the south pole goes into darkness and summer arrives in the north. Spilker points out that Cassini arrived at Saturn just a couple of years after the northern winter solstice. Scientists noticed at the time that the atmosphere in the north looked much bluer than in the south, so they will be watching to see if a similar bluish-colored haze starts to form down under as the southern hemisphere receives less sunlight.

Cassini will also keep an eye on Saturn's largest moon, Titan. The spacecraft's radar mapper has peered through Titan's smoggy skies to identify about 400 lakes, some quite large, of a liquid methane-ethane mixture. Porco describes them as being like "Lake Michigan, filled with paint thinner." The lakes are found in the polar regions, but there are more in the north than the south—including all the really big ones that might properly be called seas. Planetary scientists are curious to see if the coming of spring to the north will cause the lakes to evaporate and their contents to turn into rain in the southern hemisphere, creating fresh lakes there.

But the changing seasons alone may not be sufficient to explain the asymmetric distribution of Titan's lakes. A team led by Associate Professor of Planetary Science Oded Aharonson and including grad student Alexander Hayes (MS '08), Jonathan Lunine (MS '83, PhD '85) of the Lunar and Planetary Lab at the University of Arizona, Ralph Lorenz of the Applied Physics Lab at Johns Hopkins, Michael Allison of the NASA Goddard Institute for Space Studies, and JPL Director Charles Elachi (MS '69, PhD '71) has proposed that Titan may have a much longer cycle of climate change. Just as our ice ages are widely believed to be driven by regular, predictable variations in the tilt of Earth's axis and the eccentricity and precession of its orbit around the

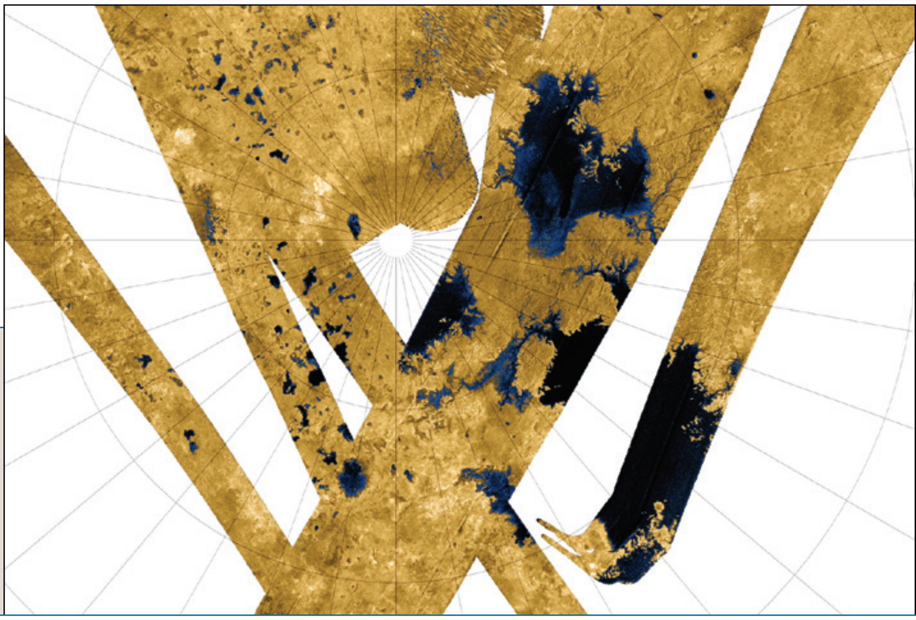

Above: The fountains of Enceladus. More than 30 jets of all sizes can be seen here, over 20 of which were previously unidentified. Cassini took this shot just before barreling through the spray on November 21, 2009, in order to sample its composition.
Right: The oily seas of Titan, tinted blue and black in this radar mosaic of the north polar region. The large one at upper-right whose coastline we can completely see is bigger than Lake Superior.

“The geysers of Enceladus,” says Porco, “are perhaps our most stupendous and significant discovery.”

sun—collectively called the Milankovitch cycles—similar changes in Saturn’s (and therefore Titan’s) tilt and orbit may be at work on Titan. If the lakes don’t fly south for the winter, it could bolster this theory.

When the time comes to end the mission, NASA may decide to send Cassini into the depths of Saturn, much as Galileo (the spacecraft, not the scientist) was incinerated in Jupiter’s atmosphere. Such a fiery demise would preclude any chance of Cassini one day crashing into—and perhaps contaminating—one of Saturn’s moons. If a later mission does discover microbes on Enceladus, for example, we’d like ironclad assurance that they’re not stowaways from Earth. First, however, Cassini would complete a final set of experiments—orbiting near Saturn’s cloud tops, perhaps sampling the dusty D ring, measuring the mass of the B ring for the first time, and acquiring precise measurements of Saturn’s gravitational and magnetic fields. In that way, Cassini will end its mission as it began, providing views of Saturn in greater detail than ever before.

By studying Saturn, says Spilker, we can put ourselves into context. Titan’s primordial chemistry—a nitrogen atmosphere, no free oxygen, and small amounts of organic materials—may show us what Earth was like a very long time ago. And Saturn’s bling may give us hints about how and where to look for planets around other stars. Notes Porco, “The processes occurring at Saturn today are similar to those that occurred in the very early days of the solar system, when the material now contained in the planets was spread out into a flattened disk. We stand to learn a great deal about the early stages of solar systems throughout the galaxy, and the cosmos, by studying Saturn’s rings.” e&s

Linda Doran is a freelance science writer, member of the National Association of Science Writers, and former Jet Propulsion Laboratory employee. Prior to that, she was a corporate communications specialist for Sandia National Laboratories in New Mexico and a reporter for the **San Gabriel Valley Tribune** ***in California. She holds undergraduate degrees in geology, anthropology, and German.***

This article was edited by Douglas L. Smith.

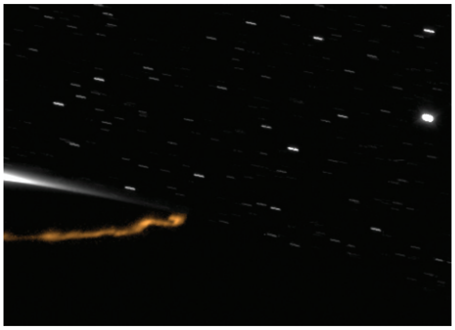

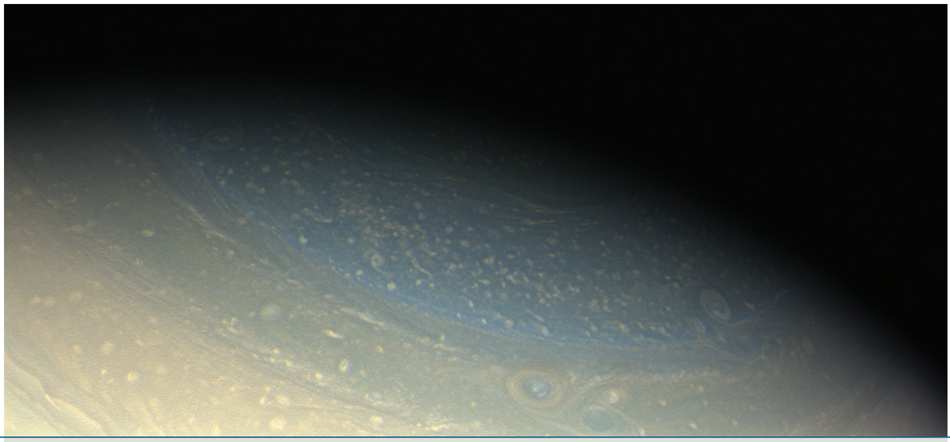

Top: A frame from the first-ever visible-light movie of Saturn’s auroras. (The aurora has been colored orange to make it easier to see.) The movie was shot over 81 hours spanning October 5–8, 2009. Bottom: Saturn’s atmosphere is not as showy as Jupiter’s, with its baroque, swirling bands of clouds in vivid colors and long-lived, eye-catching storms like the Great Red Spot. But in its own pastel way, Saturn’s atmosphere is just as complex and intriguing. This true-color image of the north pole was taken in November 2008, as the northern winter was drawing to a close. We can see hundreds of bright storm systems, and a blue that has since faded to other colors with spring’s return.

Recommend a Movie, Win a Million Bucks

Can a computer learn to choose movies you're sure to like? Artificial intelligence took on this real-world problem in the Netflix contest. A Caltech alum recounts his adventures in the quest for the million-dollar grand prize.

Many a Caltech PhD goes into research, but for the better part of a year I did so alone and unemployed—and, yes, often in my pajamas. Yet, without ever leaving my apartment, I ended up on an 11-country team chasing the million-dollar Netflix Prize in what Caltech professor Yaser Abu-Mostafa [PhD '83] has aptly described as the Super Bowl of machine learning. There were 5,169 teams competing, and with one day to go we were in first place.

Modeled after the 2003 X Prize challenge to build a privately funded spaceship, the Netflix Prize was intended to harness the creativity of computer scientists and statisticians worldwide. Humans (most of us anyway), instinctively learn from experience to recognize patterns and make predictions, but designing machines that "learn" in an even remotely similar manner will provide fodder for PhD theses for decades to come. The Netflix challenge, to build a better method for predicting a customer's future movie preferences based on past movie ratings, ultimately led to many discoveries, including profound advances in statistical modeling—as well as the fact that fans of the TV series *Friends* tend to dislike Stanley Kubrick movies.

A visitor to the Netflix website looking for something to rent faces a sea of more than 100,000 choices. Cinematch, the website's recommendation system, helps you sort through them. Cinematch bases its suggestions on a staggeringly huge database of movie ratings collected whenever one of Netflix's 10-million-plus subscribers clicks on one of the five "star" buttons below the thumbnail picture of each movie's cover art. The average customer has rated some 200 movies, so Cinematch has plenty of preference data to work with, but still, as of 2005,

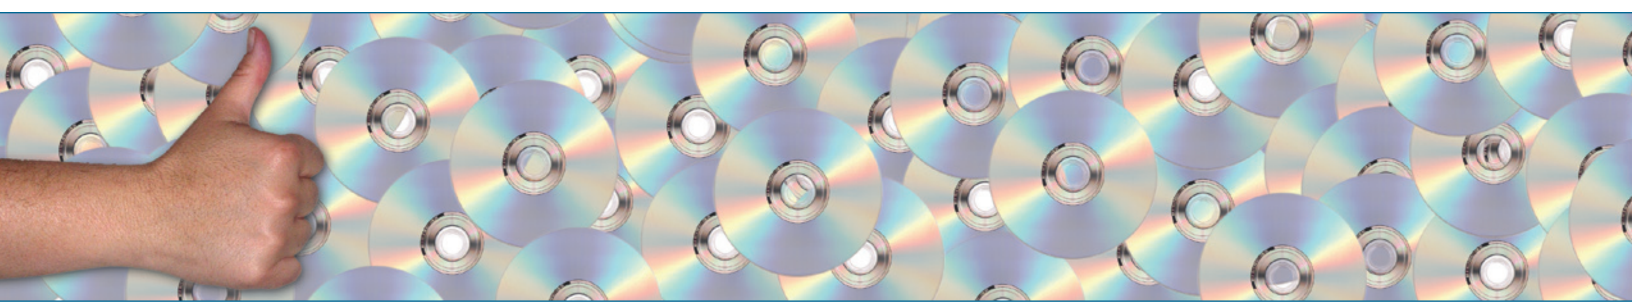

Contrary to popular belief, self-employed computer gurus do occasionally spend time in the sun. Sill enjoys a balmy spring day on Chicago's lakefront.

By Joseph Sill

Netflix was not entirely happy with it. Some of its suggestions seemed pretty random.

The company had been tweaking the system for years, and finally Netflix founder and CEO Reed Hastings tried to fix things himself. He neglected his wife and family over the 2005 Christmas vacation while manipulating spreadsheets and trying to make sense of mountains of customer data. He soon concluded that he was in over his head, and Netflix eventually decided to look outside the company for help. On October 2, 2006, Netflix formally offered one million dollars to whoever submitted the best algorithm, as long as it beat Cinematch's performance by at least 10 percent.

Entrants were allowed to download a data set of more than 100,000,000 ratings spanning seven years; 480,189 customers, represented anonymously by ID numbers; and 17,770 movies and TV series. Each data point consisted of four pieces of information: the customer ID, the movie or TV show's title, the rating, and the date the rating was made. This was the "training set"—the data you let your computer puzzle over to tease out the patterns.

Some simple calculations showed that the highest-rated movies were the *Lord of the Rings* trilogy, with the final installment, *The Return of the King*, in first place with an average rating of 4.72 stars. Interestingly, every single one of the bottom 10 movies was a horror flick. In one case, this may have been intentional, since the sixth-lowest movie (1.40 stars) was titled *The Worst Horror Movie Ever Made*. Although this film came close to living up to—or down to—its name, dead last went to *Avia Vampire Hunter* at 1.29 stars. (A zero-star rating was not permitted.) But winning the contest would require much more complex mathematics.

The competition itself involved another 2.8 million data points where only the customer IDs, titles, and dates were revealed. This data had been randomly split into two subsets, the quiz set and the test set, and contestants were not told which data points belonged to what set. Teams were allowed to upload their predictions for the entire set to the Netflix Prize server once a day. The server then compared the predictions to the true ratings and calculated the root mean squared error, or RMSE, for each set. Loosely speaking, RMSE measures the average error. If an algorithm predicted three stars when the true rating was four, for example, the RMSE would be one. The teams' RMSEs on the quiz set were posted in rank order on a public web page known as the leaderboard. The RMSEs on the test set, which would determine who—if anyone—would win the million dollars, were known only to Netflix.

Things started out swiftly, with an improvement of more than 8.5 percent over Cinematch in the first 18 months. I had a highly demanding job as a quantitative analyst for a hedge fund at the time, so I was not following the action very closely. But my old mentor, Professor of Electrical Engineering and Computer Science Abu-Mostafa, was—he was using the Netflix Prize as the basis for a project-based machine-learning course. I visited him in early 2008 to help interview prospective grad students, and he filled me in. We'd been friends for more than 15 years, so his interest excited mine; but still, my job left me no free time.

My own career in machine learning had begun in Caltech's computation and neural systems program, which occupies the intersection of computer science and engineering with neurobiology. In it, neurobiologists use computers to help understand the brain, and engineers look to neurobiology for inspiration in designing machines. With a background in applied math, I was on the theoretical periphery, and only survived the two demanding lab courses through the patience and generosity of my unlucky lab partners. One final required us to determine the inner workings of an analog VLSI chip (a particularly esoteric device) by running various diagnostic tests. I passed only by divine intervention—the fabrication plant accidentally burned the chips, and the exam was cancelled. I'm much more comfortable working with data, which in the mid-'90s meant a few thousand data points if we were lucky.

The sheer enormity of the Netflix dataset was alluring, and in the summer of 2008 a career decision provided the time to dig into it. The full fury of that fall's market

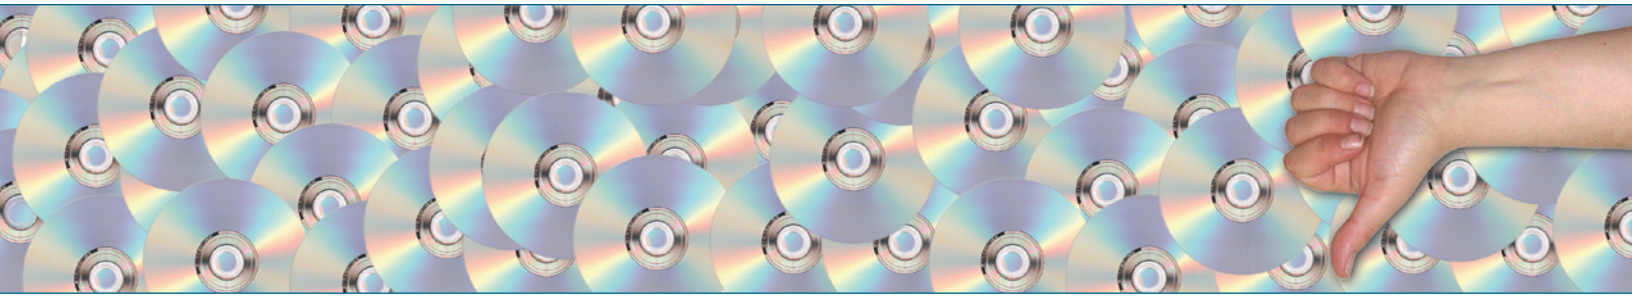

Left: A CS 156b discussion section. Abu-Mostafa (seated, at left) and grad student Panna Felson (BS '09) look on as senior Constantine (Costis) Sideris holds forth. Grad student and teaching assistant Chess Stetson is sitting on the right.

Right: At the end of 2007, BellKor ruled the Netflix leaderboard.

THE NETFLIX PRIZE AS A TEACHING TOOL

The most important feature of the Netflix competition was not the size of the prize, one million dollars, (although that didn't hurt!) but rather the size of the data: 100 million points. Machine learning researchers are used to much smaller data sets, and Netflix provided several orders of magnitude more data than what we usually have to be content with. It was a machine learning bonanza.

Machine learning lives or dies by the data that the algorithm learns from. The whole idea of this technology is to infer the rules governing some underlying process—in this case, how people decide whether they like or dislike a movie—from a sample of data generated by that process. The data are our only window into the process, and if the data are not enough, nothing can be done but make guesses. That doesn't work very well.

I saw the Netflix data as an opportunity to get my CS 156 Learning Systems students to try out the algorithms that they had learned in class. With 100 million data points, the students could experiment with all kinds of ideas to their heart's content.

The first time I gave the Netflix problem as a class project, the competition was still going on. For the project, each team had to come up with its own algorithm and maximize its performance. Then all the algorithms would be blended to give a solution for the class as a whole. Like Joe, our hope was to rapidly climb the leaderboard with each submission.

Perhaps my biggest challenge as an instructor was to work out a way of ensuring that the various teams tried different techniques. As you will see, blending radically different solutions is key to getting good performance, so if everybody tried the same approach because it seemed to be the most promising technique known, blending the results of such duplicative efforts would not give a lot of improvement.

Therefore, I announced that a team's grade would not depend on their algorithm's individual performance, but on the incremental improvement in the class-wide solution when the algorithm was incorporated into the blend. This gave the students an incentive to explore the less-traveled roads that might offer a better chance of shining in the blended solution. There was great educational value in pushing people to venture "outside the box."

How did we do? Well, I registered two team names with Netflix. I told the students that if we did really, really well, we would submit under "Caltech." The other name was that of a rival school back east, and we would use it if we did really, really badly. It turned out that we fared neither too badly nor too well. Our effort gave about a 6 percent improvement over the original Cinematch system, so we did not officially submit it.

The second time the course was offered, the Netflix competition had already ended. There was no blending this time around, and the grading process was more of a judgment call.

With the benefit of hindsight, the class' performance was much better. One team was getting a weekly improvement that it took the teams in the actual competition months to achieve. In fact, at one point a guy in the class wanted to set an appointment to meet with me. He had the top individual score of 5.3 percent, so I jokingly told him that he needed to get to 6 percent before I would see him. Within a few days he had gotten to 6.1 percent. —*YA-M* E&S

meltdown had not yet hit, and my employer was doing fine, but I was ready to move on from finance. I needed some time off, and I needed a project to keep my skills sharp while I mulled things over. The Netflix Prize was perfect.

GETTING UP TO SPEED

The contest had now been going on for almost two years, so I had a lot of catching up to do. But I had ample time, now that I had quit my job, and there was a road map to the top of the leaderboard. To keep things moving, Netflix was offering an annual $50,000 "Progress Prize" to the team with the lowest RMSE, as long as it represented at least a 1 percent improvement on the previous year's best score. To claim the prize, however, you had to publish a paper describing your techniques at a reproducible level of detail. The first Progress Prize had gone to BellKor, a team from AT&T Research made up of Yehuda Koren (now with Yahoo! Research), Bob Bell, and Chris Volinsky. All I had to do was follow their instructions, and in theory I should get similar results.

As I read BellKor's paper, I quickly realized that ironic quotation marks belonged around the phrase "all I had to do." Their solution was a blend of over 100 mathematical models, some fairly simple and others complex and subtle. Fortunately, the paper suggested that a smaller set of models—perhaps as few as a dozen—might suffice. Even so, nailing down the details would be no easy task.

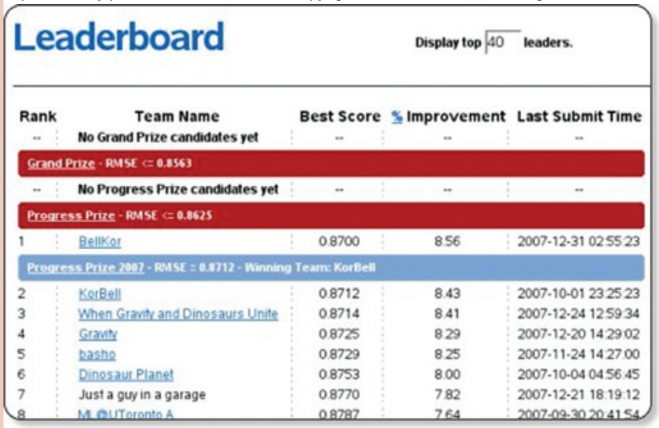

Several other leading teams had also published papers, even though they were not compelled to do so, and some of their methods also looked promising. Three general approaches seemed to be the most successful: nearest neighbors, matrix factorization, and restricted Boltzmann machines.

Nearest-neighbor models are among the oldest, tried-and-true approaches in machine learning, so it wasn't surprising to see them pop up here. There are two basic variations—the user-based approach and the product-based approach.

Suppose the task is to predict how many stars customer 317459 would give *Titanic*. The user-based approach would look for like-minded people and see what they thought. This similarity is measured by calculating correlation coefficients, which track the tendency of two variables to rise and fall in tandem. However, this approach has some serious drawbacks. The correlations between every possible pairing of customers need to be determined, which for Netflix amounts to 125 billion calculations—a very heavy computational burden. Furthermore, e-commerce companies have found that people tend to trust recommendations more if the reasoning behind them can be explained. "Selling" a user-based nearest-neighbor recommendation is not so easy: "There's this woman in Idaho who usually agrees with you. She liked *Titanic*, so we think you will, too—trust us on this one."

For these reasons, product-based nearest-neighbor techniques are now much more common, and the Netflix contest was no exception. With 17,770 movies, as opposed to half a million customers, computing the correlations between all possible pairings in the training set was much more manageable, and the resulting table fit easily into a computer's RAM. These precomputed correlations led to odd insights—it's how I learned that fans of *Friends* tend to dislike Stanley Kubrick. (If you correlate Netflix's catalog with the ratings given to a DVD of a season of *Friends* and rank the results, *Dr. Strangelove* and *2001: A Space Odyssey* end up at the bottom.) Typically, only positive correlations are used—for instance, *Pulp Fiction* is highly correlated with other Quentin Tarantino movies like *Reservoir Dogs* and *Kill Bill,* as well as movies with a similar dark and twisted sensibility like *Fight Club* and *American Beauty*. Thus a product-based approach allows Netflix to say, "since you enjoyed *Fight Club*, we thought you might like *Pulp Fiction*."

Surprisingly, I found that negative product-based correlations were nearly as useful as positive ones. Although the connections were not as obvious, there was often a certain logic to them. For instance, *Pulp Fiction* and Jennifer Lopez romantic comedies like *The Wedding Planner* and *Maid in Manhattan* were strongly anticorrelated, and fans of the TV series *Home Improvement* generally disliked quirky movies such as *The Royal Tenenbaums*, *Eternal Sunshine of the Spotless Mind*, and *Being John Malkovich*. I ended up devising a negative-nearest-neighbors algorithm—a "furthest opposites" technique, if you will—that scored an RMSE of 0.9570, which was nearly as good as the 0.9513 of the original Cinematch software. I got perverse enjoyment out of imagining a customer being told that since he hated *Rambo III*, he might like *Annie Hall*.

Nearest-neighbor systems improved greatly over the contest's first year, as teams delved into the math behind the models. These new twists represented significant progress in the design of recommendation systems, but a bigger breakthrough came in the form of another technique called matrix factorization.

Matrix factorization catapulted into prominence when "Simon Funk" suddenly appeared out of nowhere, vaulting into fourth place on the leaderboard. Funk, whose real name is Brandyn Webb, had graduated from UC San Diego at 18 with a computer science degree, and has since had a spectacular career designing algorithms. He was quite open about his method, describing it in detail on his website, and matrix factorization eventually became the leading technique within the Netflix Prize community.

Matrix factorization represents each customer and each movie as a vector, that is, a set of numbers called factors. The customer's predicted rating of that movie is the dot product of the two vectors—a simple mathematical operation involving multiplying each customer factor by each movie factor and summing up the result. Essentially, the movie vector encodes an assortment of traits, with

DOT PRODUCTS

Let's say that movie vectors have four factors: sex (S), violence (V), humor (H), and music (M). Maggie, a customer who has just watched *Brigadoon* and is looking for more of bonnie Scotland, has preferences S = 2, V = 1, H = 3, M = 5. If *Braveheart* scores S = 2, V = 5, H = 1, M = 1 on this scale, then the dot product predicting Maggie's likely reaction to Mel Gibson's gorefest is $(2 \times 2) + (1 \times 5) + (3 \times 1) + (5 \times 1) = 17$. Since the maximum possible score in this example is 100, she should probably give *Braveheart* a miss.

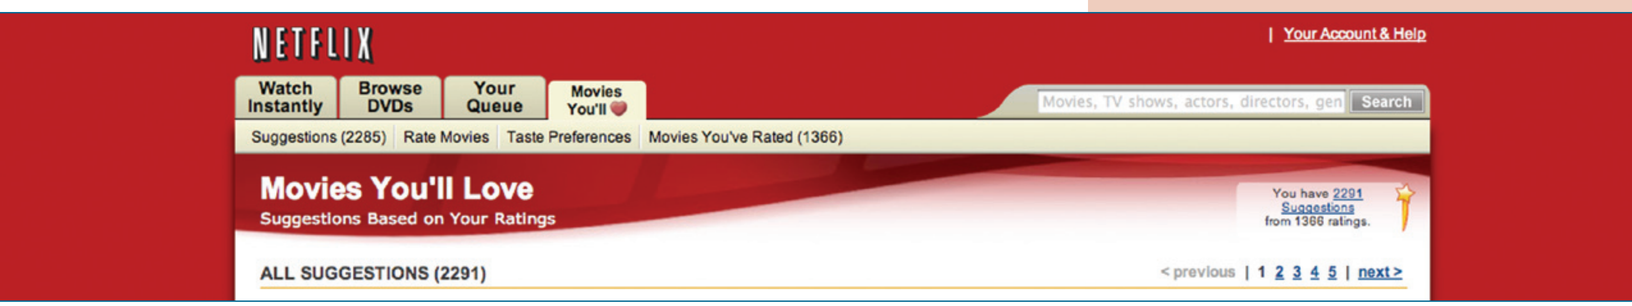

each factor's numerical value indicating how much of that trait the movie has. The customer vector represents the viewer's preferences for those traits. It's tempting to define the factors ahead of time—how much money the film grossed, how many Oscar winners are in the cast, how much nudity or violence it has—but in fact they are not predetermined in any way. They are generated by the model itself, inside the "black box," and only the model knows exactly what they mean.

The default number of teams displayed on the web page was 40, so if I cracked the Top 40 I could at least point myself out to people easily.

All we know is that as the model learns, it continually adjusts the factors' values until they reliably give the right output, or rating, for any set of inputs.

But if a model had just a handful of factors, sometimes their meanings leapt out. One such factor tagged teen movies: the *American Pie* series and *Dude, Where's My Car?* came out on top, and old-school classics like *Citizen Kane* and *The Bridge on the River Kwai* were at the bottom. Another loved romantic comedies such as *Sleepless in Seattle* and hated the *Star Trek* franchise. A third chose left-leaning films, placing Michael Moore movies such as *Fahrenheit 9/11* highest and movies such as *Celsius 41.11* (subtitled *"the temperature at which the brain begins to die"* and intended as a Republican response) lowest. As the number of factors grew—beyond 100, in some cases!—they got much harder to interpret. The trends being discerned got more subtle, but no less real.

The third class of approaches, the restricted Boltzmann machine, is hard to describe concisely. Basically, it's a type of neural network, which is a mathematical construct loosely inspired by the interconnections of the neurons in the brain. As a computation and neural systems alumnus, I was glad to see neural networks well represented.

BellKor's $50,000 winning blend incorporated numerous twists and tweaks of all three major techniques as well as several less-prominent methods. This blending approach became standard practice in the Netflix Prize community. The process of blending predictions was itself a meta-problem in machine learning, in that the blending algorithm had to be trained how to combine the outputs of the component models.

Teams working with blends soon discovered that it didn't make sense to fixate on improving any particular technique. A better strategy was to come up with new models that captured subtle effects that had eluded your other models. My favorite example was a model based solely on what day of the week the rating had been made. Even after controlling for the types of movies that get rated on different days of the week, movies were more likely to get a bad rating on Mondays and more likely to get a good rating on weekends. On its own, this model performed horribly, but if none of the other models took the day of the week into account, this tiny but statistically significant signal could lead to a noticeable boost in the ensemble's accuracy.

BLENDED MODELS, BLENDED TEAMS

By the fall of 2008, having familiarized myself with all the major techniques and the art of blending large numbers of models, I set out to climb the leaderboard. I started with what I thought was a modest goal—the default number of teams displayed on the web page was 40, so if I could crack the Top 40 I could at least point myself out to people easily. Although my girlfriend was being very supportive, I wanted a tangible achievement to show for my solitary efforts hunched over the desk in the living room of my Chicago apartment, pounding away on my laptop. But with more than 5,000 teams competing, landing in the Top 40 meant being in the 99th percentile. I should have realized it wouldn't be so easy.

For one thing, there were subtleties unmentioned in the papers that made the difference between a working model and a failure. Deducing these tricks on my own sometimes took weeks at a time. Another challenge was weighing how much time to spend on known methods versus inventing original approaches. I probably spent too much time trying for a home run—a stupendous, brand-new technique.

Still, I scored one triumph. When blending their models, most of my competitors relied on linear regression, a time-honored technique whose many virtues include a straightforward, analytically exact procedure for obtaining the best fit to the data. My intuition suggested that the blend should be adaptive, depending on such side information as the number of ratings associated with each customer or movie. By making every model's coefficients a linear function

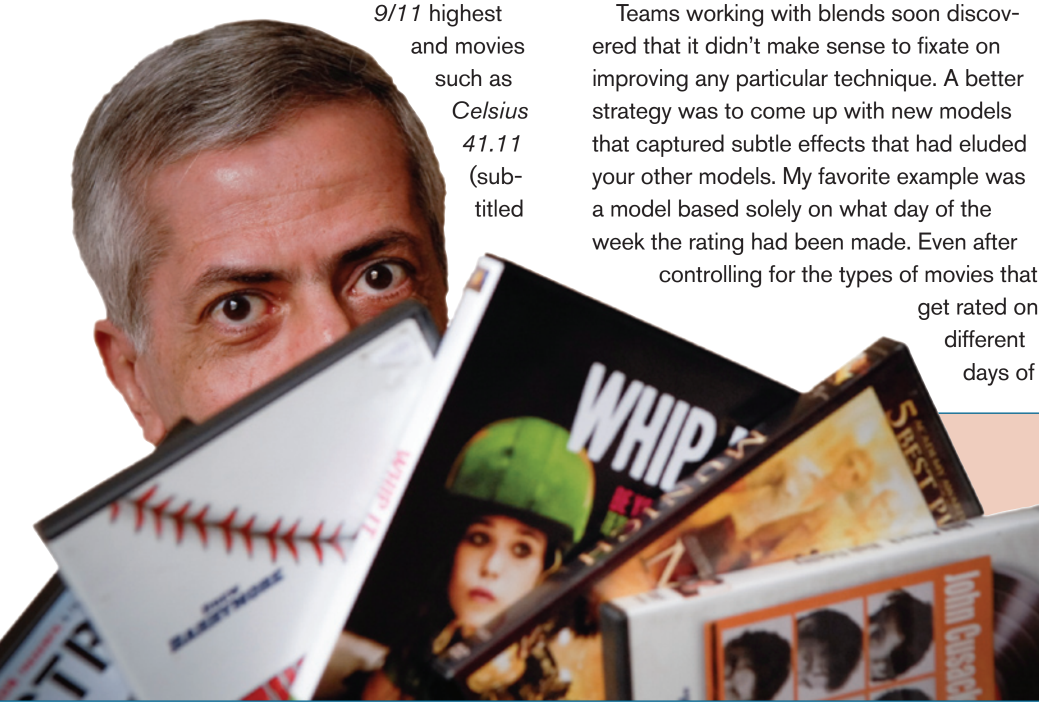

Abu-Mostafa uses telepathy to pick movies for you. Machine learning systems can't do that yet.

Time was running out in the spring of 2009, but nobody knew it. Then suddenly, on June 26, a newly formed team reached the goal. The competition promptly shifted into overdrive, as the contest included a one-month "last call" for all comers to take (or defend) the lead. The team with the best score at the end of that month would be declared the winner.

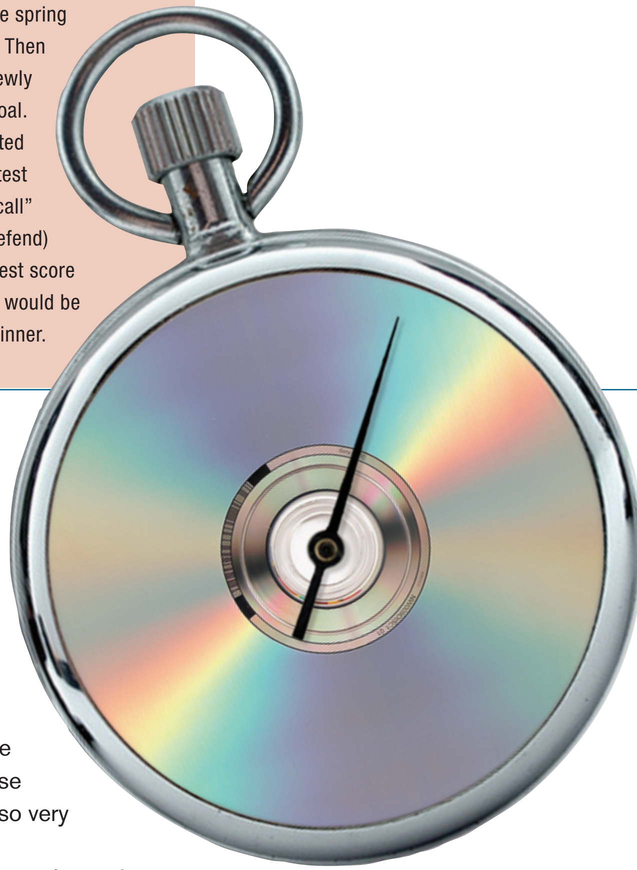

of these side variables, I got a significant boost in accuracy while still retaining linear regression's key merits.

Despite this breakthrough, come January 2009 I was still hovering just below where I wanted to be, stuck in the mid-40s. I would often peruse the top 10 with a mixture of respect and jealousy. One day, a new team suddenly appeared there—the Grand Prize Team, a name that struck me as presumptuous and cocky. When I read about them, though, I was intrigued. GPT's founders—a team of Hungarian researchers called Gravity, and a team of Princeton undergrads named Dinosaur Planet—had merged in 2007. That team, dubbed When Gravity and Dinosaurs Unite, had very nearly won the first Progress Prize, but was overtaken by BellKor in the final hours.

GPT issued a standing invitation to all comers to join them, but in order to be admitted the applicant had to demonstrably improve GPT's score—not an easy task. As of GPT's founding, their RMSE stood at 0.8655. The million-dollar goal was 0.8563, so only 0.0092—or 92 basis points, as they were called—remained to go.

Shaving off just one basis point was excruciatingly hard, and GPT's offer reflected this. The original members would only claim one-third of the prize, or $333,333, and the remainder would be split among the new members in proportion to their basis-point contribution. Thus one basis point, the smallest measurable improvement, was worth almost $7,000. This may seem overly generous, but it reflected how difficult progress had become. Weeks or even months would go by with nothing happening at the top of the leaderboard, and a boost of just a few basis points was cause for celebration. The old 80/20 rule—that 80 percent of the payoff comes in the first 20 percent of the work—was in full force.

My advancement as a lone wolf stymied, I crossed my fingers and sent my adaptive-blending code to GPT captain Gábor Takács. Reading the email I got in response was the most satisfying moment I had yet experienced—I'd boosted their score by more than 10 basis points! GPT's next official submission to Netflix a few days later jumped from 0.8626 to 0.8613 on the strength of my contribution, leapfrogging a few other entrants in the process. As of February 2009, I was suddenly a significant shareholder on a leading team.

However, first place belonged to a merger of BellKor and two Austrian graduate students calling themselves Big Chaos. The amalgam, BellKor in Big Chaos, had won the second Progress Prize in October 2008, with a score of 0.8616. They had since climbed to 0.8598, and there they had ground to a halt. It seemed like that figure was chiseled in stone—so close to the magic number, and yet so very far away.

My stake as one of a dozen members of GPT gave me newfound motivation, as did the contrast between our leap and our competitors' glacial progress. We were still underdogs, but underdogs with momentum on our side. Over the next few months, I found some other side variables I could exploit, and we crept up the leaderboard. A new recruit from Israel, Dan Nabutovsky, boosted our score by another six basis points.

Our collaborative style varied. Some GPT members (or prospective members) simply sent Gábor their predictions. He'd add them to the mix and see whether the overall blend improved—something he could do without even knowing how the new models worked. Gábor and I, however, were collaborating at a much deeper level, exchanging code and debating which approaches were most likely to give us a boost.

Meanwhile, another underdog was also making a run. Calling themselves Pragmatic Theory, Martin Chabbert and Martin Piotte of Quebec had been steadily rising up the leaderboard. Neither of them had any formal training in machine learning, and both were holding down full-time jobs. Nonetheless, by March 2009 they had surpassed BellKor in Big Chaos and taken the lead with a score of 0.8597. After a bit of jockeying, the leaderboard settled into a new equilibrium, with the two teams tied at 0.8596. As the weeks dragged on, GPT narrowed the gap, but the front-runners continued to inch ahead. I'd check the standings every day, and often several times a day, wincing on the rare days when a new score was posted. Their progress was so slow that I felt we had plenty of time. I was wrong.

ENDGAME

Late afternoon on Friday, June 26, I checked the leaderboard, and what I saw felt like a punch in the gut. A new team—BellKor's Pragmatic Chaos—had broken the million-dollar barrier, with a score of 0.8558. They hadn't won yet, though. The contest rules provided for a one-month "last call" during which everyone had a final shot at the prize, and the team with the best score at the end of this period would be declared the winner. So we weren't toast, but things didn't look good. We were 36 basis points behind, a gap that felt as wide as the Grand Canyon. We'd been running an ultramarathon—months-long for me; years-long for some of my teammates—and only the day before, we'd been on the heels of the leaders. Now, just short of the finish line, we looked up and found they were so far ahead that we had to squint to see them.

We decloaked on Saturday afternoon. With less than 24 hours to go, we were on top.

With so little time left, GPT became much more collaborative, with most members exchanging detailed information about their algorithms via email and chat sessions. Everyone else was scrambling to catch up as well. A few leaders were on a quixotic quest to win on their own, but many could see that there was strength in numbers. During the first two weeks of the final month, another coalition formed. Vandelay Industries, the name of a fake company from an episode of *Seinfeld*, and Opera Solutions, a real consulting firm, united and began aggressively recruiting anyone in the top 100. Soon their score was nearly the equal of ours, but even so, we were both still far behind. It was clear to all what we had to do. We negotiated a merger and a 50/50 split of the prize money, halving the value of my basis points. We were now a group of 30 people, and since the technique of blending models is known in the trade as ensemble learning, we named ourselves The Ensemble.

We worked hard but in stealth, keeping our new name off the leaderboard. BellKor's Pragmatic Chaos might have tried to recruit the uncommitted if they felt threatened, so we hoped to lull them into complacency by keeping our existence a secret. We had developed accurate ways to estimate our score in-house, and before long we determined that we could surpass the 0.8563 million-dollar barrier. Meanwhile, the opposition also inched ahead, to 0.8555. Little did they know that their lead was shrinking.

Our efforts intensified as the contest entered its final week. With team members in Europe, India, China, Australia, and the United States, we were working the problem 24/7. We now had thousands of models, and there were countless ways to blend them. We could even blend a collection of blends. In fact, our best solution was many-layered—a blend of blends of blends of blends. If we won, we'd have to document our methods before being awarded the prize, and I began to worry whether we'd be able to reconstruct what we had done. With so little time left, though, we decided to deal with that later.

The competition was scheduled to end on Sunday, July 26, 2009, at 1:42 p.m. Chicago time. That Thursday, we obtained a thrilling result: 0.8554. We could take the lead! Ah, but should we do so publicly? As we phrased it in our email discussions, should we decloak? Our members' agendas varied—some of us were all about winning, while others were more interested in the publicity. If we decloaked too soon, we might spur BellKor's Pragmatic Chaos to work extra hard or go on a recruiting binge. On the other hand, our window of opportunity might not last. If our next submission was merely the runner-up, we'd get less media attention than if we had suddenly seized first place. Ultimately, we decloaked on Saturday afternoon. With less than 24 hours to go, we were on top.

However, we were Number One with an asterisk. Remember that the leaderboard posted the quiz set's results, while the true winner would be the highest performer on the test set—a ranking known only to the Netflix engineers running the contest. Since both sets were drawn from the same data, the two scores should be close, but precisely how close? Our learning and tweaking process had been influenced in subtle ways by feedback from the quiz set—an example of what's known in the education biz as "teaching to the test." What would happen now that the test itself had changed?

Our first-place debut brought a rush of adrenaline, but the game was far from over. We planned to make one last submission a few minutes before the contest ended. We worked frantically through Saturday night and into Sunday morning. With 20 minutes to go, BellKor's Pragmatic Chaos matched our quiz score of 0.8554. A cacophony of emails and chat messages ensued, and as the final minutes ticked away, Peng Zhou announced from Shanghai that he thought he could achieve 0.8553. We sent in his blend, retook the lead, and the deadline passed. We had finished first, but had we won?

Shortly after the final buzzer sounded, Netflix announced on the Netflix Prize discussion board that two teams had qualified, and that their submissions were being evaluated. We still didn't know who'd won, but at least we knew we'd achieved a 10 percent (or better!) improvement on the test set. We were relieved that most of our accuracy on the quiz set had carried over.

In previous years, the Progress Prize winners had been notified via email shortly after the deadline but well before any public announcement. We had a gentleman's agreement with our rivals that if either team received such an email, we'd notify the other. An agonizing 90 minutes passed, and finally, around 3:00 p.m., we got the email we dreaded. BellKor's Pragmatic Chaos had won.

I went for a long jog to clear my head. I was a little depressed, but I wasn't feeling so bad. It was safe to say that more of their original work was in our solution than vice versa, since they had been required to publish two papers in order to win the Progress Prizes. I was also in awe of Pragmatic Theory, who had pulled off a stunning victory in a field where they were newcomers working in their spare time. I couldn't begrudge them the victory.

Netflix's official stance for the next two months was that the contest had not yet been decided, as the winning software was still being validated. This led to a fair amount of confusion. We still held first place on the leaderboard, and it was entirely reasonable for a casual observer to assume that we had won. Netflix asked us to say only that we were happy to have qualified, which led to some awkward situations.

Netflix finally announced the winner at a press conference in New York on September 21. It was there that we learned that we had, in fact, tied, with both teams scoring an RMSE of 0.8567 on the test set. However, BellKor's Pragmatic Chaos had sent in their submission 20 minutes before we did. The pain of losing by such a minuscule margin was somewhat allayed by being quoted in the *New York Times*, and I bought several copies as souvenirs.

The press conference was my first opportunity to meet both my teammates and my competitors. I asked the Pragmatic Theory duo whether they were ready to quit their jobs and become machine-learning researchers. I was only half joking, but Martin Piotte just shrugged and said he might take up another hobby instead. After less than 18 months in the field, and having published perhaps the year's most important paper, I guess he felt it was time to move on.

At the awards ceremony that followed, some members of BellKor's Pragmatic Chaos gave a talk outlining the advances made in their final push. It turned out that Chabbert and Piotte had used the dates on which the ratings had been made in a very creative way. A movie could be rated in one of two contexts: either within a few days of watching it (Netflix asks customers to rate the movies they've just rented), or when browsing the website and rating previously seen movies. If a customer rated dozens or hundreds of movies on a single day, it was a safe bet that most of those films had been seen a long time ago. However, certain movies got better ratings immediately after viewing, while others fared better in retrospect. I had done a little work along these lines myself, and had squeezed out an additional 1-basis-point share in GPT as a result, but evidently the effect was more significant than I realized. Pragmatic Theory modeled it in detail and got a significant accuracy boost as a consequence.

Another interesting advance came from Big Chaos's Austrians. Most contestants optimized each individual model's accuracy on its own, and then blended it into the collection and crossed their fingers, hoping that the overall prediction improved. Big Chaos had found a way to train the individual models to optimize their contributions to the blend, rather than optimizing their own accuracy.

The power of blending, which the contest demonstrated over and over again, was the Netflix Prize's take-away lesson. You could even reverse-engineer a blend of models, using the results to design an "integrated" single model that incorporated the effects captured by the various simpler models. Indeed, Pragmatic Theory created a single model that scored 0.8713, equal to the RMSE of the 100-model blend that had won BellKor the first Progress Prize.

Although not all of the techniques involved in the winning solution have been incorporated into the working version of Cinematch, many of them have, and Netflix has seen increased customer loyalty since implementing these advances. Hastings, the Netflix CEO, said in the *New York Times* that the contest had been "a big winner" for the company. As for myself, I didn't quite finish as a winner in the formal sense, but I'm still thrilled with how the experience turned out. The contest has led to actual paid consulting work, and a group of my teammates and I are writing a paper on some of our techniques. Many, many other papers have and will come out of the contest, to the great benefit of the broader machine-learning research community. Only one team won the million dollars, but the Netflix competition ended up producing prizes of many kinds. E&S

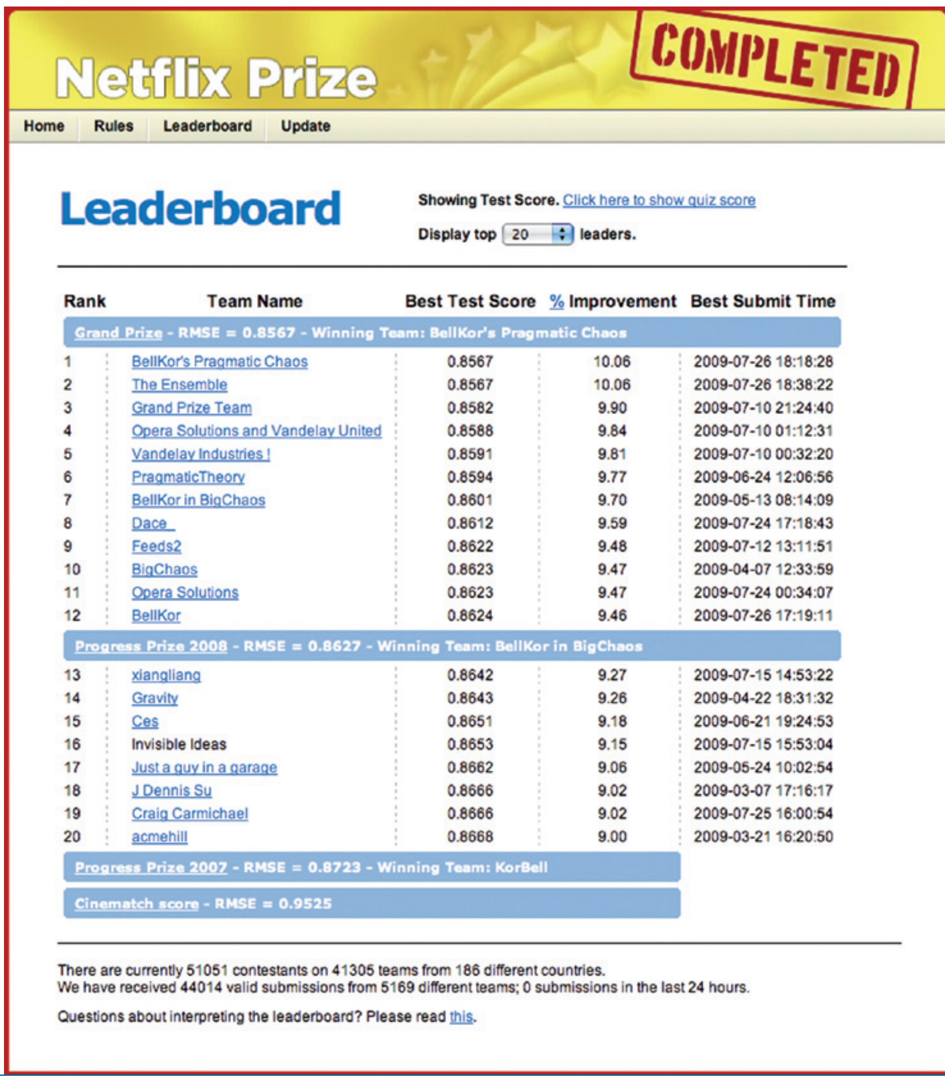

Leaderboard

Showing Test Score. Click here to show quiz score

Display top 20 leaders.

Rank	Team Name	Best Test Score	% Improvement	Best Submit Time
Grand Prize - RMSE = 0.8567 - Winning Team: BellKor's Pragmatic Chaos				
1	BellKor's Pragmatic Chaos	0.8567	10.06	2009-07-26 18:18:28
2	The Ensemble	0.8567	10.06	2009-07-26 18:38:22
3	Grand Prize Team	0.8582	9.90	2009-07-10 21:24:40
4	Opera Solutions and Vandelay United	0.8588	9.84	2009-07-10 01:12:31
5	Vandelay Industries !	0.8591	9.81	2009-07-10 00:32:20
6	PragmaticTheory	0.8594	9.77	2009-06-24 12:06:56
7	BellKor in BigChaos	0.8601	9.70	2009-05-13 08:14:09
8	Dace_	0.8612	9.59	2009-07-24 17:18:43
9	Feeds2	0.8622	9.48	2009-07-12 13:11:51
10	BigChaos	0.8623	9.47	2009-04-07 12:33:59
11	Opera Solutions	0.8623	9.47	2009-07-24 00:34:07
12	BellKor	0.8624	9.46	2009-07-26 17:19:11
Progress Prize 2008 - RMSE = 0.8627 - Winning Team: BellKor in BigChaos				
13	xiangliang	0.8642	9.27	2009-07-15 14:53:22
14	Gravity	0.8643	9.26	2009-04-22 18:31:32
15	Ces	0.8651	9.18	2009-06-21 19:24:53
16	Invisible Ideas	0.8653	9.15	2009-07-15 15:53:04
17	Just a guy in a garage	0.8662	9.06	2009-05-24 10:02:54
18	J Dennis Su	0.8666	9.02	2009-03-07 17:16:17
19	Craig Carmichael	0.8666	9.02	2009-07-25 16:00:54
20	acmehill	0.8668	9.00	2009-03-21 16:20:50
Progress Prize 2007 - RMSE = 0.8723 - Winning Team: KorBell				
Cinematch score - RMSE = 0.9525				

There are currently 51051 contestants on 41305 teams from 186 different countries.
We have received 44014 valid submissions from 5169 different teams; 0 submissions in the last 24 hours.

Questions about interpreting the leaderboard? Please read this.

Joseph Sill is an analytics consultant. He earned his Caltech PhD in 1998, and a BS in applied math from Yale in 1993, where he was elected to Phi Beta Kappa. Before competing for the Netflix Prize, he had worked for Citadel Investment Group and NASA's Ames Research Center.

This article was edited by Douglas L. Smith.

It ain't over till it's over, as Yogi Berra used to say. Well, it's finally over, and 20 minutes made all the difference.

A TRUE ALTERNATIVE FOR YOUR FAMILY

We're not your run-of-the-mill, try-to-do-everything Credit Union. We choose to deliver only those products and services that we can do BEST. And year after year, this selective philosophy works!

The result is an uncommonly successful and cost-efficient operation that is able to return a far greater percentage of its earnings to its members than the industry average.

We consistently offer among the highest savings rates in the nation...the lowest loan rates available...and state-of-the-art online access. All with the unsurpassed safety and warm, personal service that has made Caltech Employees Federal Credit Union the trusted choice of our members — for both saving and borrowing — for over half a century.

As a proud member of the extended Caltech family, we never forget that we work for YOU!

So, when we offer a financial product, you can be sure it will be one of the best. That's our policy, and the heart of our vision: to always be YOUR **FIRST CHOICE** FOR FINANCIAL SERVICES.

Expect more from the Credit Union exclusively serving the extended Caltech & JPL community.

CALTECH EMPLOYEES Federal Credit Union

800.592.3328 • www.cefcu.org

A Strong, Secure and Exclusive Financial Resource

Foothill Office: 528 Foothill Blvd. • La Cañada Flintridge
JPL Office: 4800 Oak Grove Dr., Bldg. 218 • Pasadena
Campus Office: 515 S. Wilson Ave. • Pasadena

Must qualify for CEFCU membership to join. Minimum $5 deposit and 25¢ membership fee due upon opening any CEFCU account. All loans subject to credit approval. CEFCU is an Equal Opportunity Lender.

Your savings federally insured to at least $250,000 and backed by the full faith and credit of the United States Government
NCUA
National Credit Union Administration, a U.S. Government Agency

Engineering & Science welcomes letters. Send correspondence to Douglas L. Smith, editor, *E&S* magazine, Caltech mail code 1-71, Pasadena, CA 91125, or email dsmith@caltech.edu. We reserve the right to edit any letters selected for publication for length, content, and clarity.

LETTERS

SUNLIGHT IN YOUR TANK

First a general comment: the Fall 2009 issue of *E&S* seems meatier than usual. Thank you.

While the science from Professor Sossina Haile's group is interesting, isn't the pursuit of chemical fuels from sunlight for use in internal combustion (or even gas-fired turbine) engines a bit of a futile exercise if the energy conversion in the engine is only 25%? 25% × 25% = 6% overall (sunlight to fuel × fuel to mechanical). Photovoltaic and electric motors seem to do better at 13% × 80% = 10% (sunlight to electrical × electrical to mechanical). Even if one adds storage, for example 70% for pumped hydroelectric, the final result is 7% versus 6%.

Vertically migrating zooplankton would seem to contribute to vertical transport of energy in at least two other ways—the current needed to maintain position in the water column if they are not neutrally buoyant, and the transfer of heat when they descend from warmer surface waters coupled with the "heat deficit" when they return to the surface after spending time at depth.

Phelps Freeborn [BS '65]

▪ ▪ ▪

In his letter to *E&S* in the Winter 2010 issue, Frank Weigert (PhD '68) comments on the brilliant research reported in the previous issue that gives some hope of finding a solution for CO_2 recycling.

He noted that the process requires high temperature from a solar power tower to strip the catalyst of its oxygen, then using the catalyst to strip oxygen out of a mixture of steam and CO_2 to produce mixture of H_2 and CO known as syngas.

The syngas can then be used as feed stock for a Fisher-Tropsch unit to produce usable fuels.

Weigert quite rightly makes the point that solar power towers and FT processes, as they're known today, are not cheap and cannot compete with oil and gas production.

He forgets that science and engineering never stands still, as proven all the time in *E&S* magazine.

I can't comment on the future cost of solar tower power, but I can on FT processes.

The Oxford Catalysts Group and its subsidiary, Velocys Inc., has invented a microchannel reactor and a hyperactive FT catalyst. The two of them together outperform the best in class fixed-bed reactor by a factor of 15.

This opens the possibility of economic applications of small-scale gas-to-liquid or waste-to-liquid. One demonstration unit is being installed in Güssing, Austria, to test a feedstock from a wood-chip gasifier, and in Brazil to test its potential to avoid gas flaring in offshore production.

In other words, science never stands still, and every avenue to stop or to recycle CO_2 should be investigated with an open mind.

Pierre Jungels [PhD '73]

THE ARROW OF TIME

Sean Carroll, in his interesting article, "The Arrow of Time," says that "effects always follow causes. It turns out that *all* (emphasis in the original) of these phenomena can be traced back to the second law [of thermodynamics]." I am a big fan of entropy—I've even written poems about it—however, there are examples of effects following causes that are unrelated to changes in entropy.

Consider the following closed system. One atom of carbon-13 is in its first excited state and stationary in the rest frame of the observer, and another atom of carbon-13 is in its ground state and moving with respect to the observer along a vector directly toward the first atom, and at a speed equal to the recoil speed that the first atom will have as it makes a radiative transition to the ground state. Assume further that the photon emitted by the first atom is directed at the second atom and is absorbed, putting the second atom in its first excited state.

Causality requires that the absorption of the photon occurs at a later time than the emission of the photon. This is certainly true in the observer's rest frame, and according to Einstein, it will be true in all frames. Please correct me if I'm wrong about that. The final state is not the same as the initial state, but it is essentially the same as a time-reversal of the initial state, so it will have the same entropy. I do not think that quantum mechanical uncertainties—line width of the ground and excited states, the fact that the atoms' positions and velocities have uncertainties, and that the photon energy is also somewhat uncertain, since it has a finite lifetime—will change the conclusion. (Again, please correct me if I'm wrong).

Not only must the final state occur at a later time than the initial state, there is the intermediate state—consisting of two atoms, both in the ground state, and a photon—which must occur after the initial state and before the final state, and which will have the same entropy.

Bill Tivol [BS '62]

▪ ▪ ▪

The article in the Winter 2010 issue of *E&S* entitled "The Arrow of Time" had some interesting information, but neglected to mention two very important applications of the second law of thermodynamics in understanding today's crisis in energy and mineral resources.

Any machine produces work (mechanical energy) by transferring heat (thermal energy) from a high-temperature reservoir, usually some sort of energy-consuming heat source, to a low-temperature reservoir, or heat sink; the second law of thermodynamics gives the maximum efficiency of this process as $(T_{hi} - T_{lo})/T_{hi}$, where T_{hi} is the temperature of the heat source and T_{lo} is the temperature of the heat sink. In order to achieve an efficiency of one (the holy grail of all perpetual motion machines), the engine would have to exhaust into a reservoir having a temperature of absolute zero.

Naive optimists viewing the ever-increasing cost of mineral resources, particularly fossil fuels, are fond of saying that "we can always find more by digging deeper." Unfortunately, the second law of thermodynamics tells us that our exploitation of the limited number of highly concentrated mineral resources has resulted in a vast increase in entropy as these resources are chemically transformed and/or spread around the world. This is a process that can't be reversed, just like Sean Carroll's arrow of time. It is obvious that the key to the survival of our civilization in the near term, i.e. the next several hundred years, is to switch to more renewable energy resources and to recycle precious materials such as gold. Digging deeper is not the answer; the high concentrations of minerals are limited to the earth's crust, which is less than 30 miles thick).

Of course, it is possible that our planet could cycle through another period of several hundred million years of high temperatures and enough carbon dioxide to produce a biomass large enough to lay down another rich layer of fossil fuels. Or the planet could undergo extreme volcanic activity to produce new mineral deposits by bringing precious elements to the surface in very localized magmatic intrusions. Sadly, this would take a very long time, during which the planet would not be very habitable for *homo sapiens*.

Peter Gottleib [BS '56]

▪ ▪ ▪

I just read "The Arrow of Time" in the latest *E&S*. It sounds like the author thinks that it is possible to know the exact state of the universe, but the Heisenberg uncertainty principle states that this is not possible. In addition, Conway and Kochen's free will theorem denies any possibility of "hidden variables"—assuming that humans have free will; this cannot be proven, but without it, all human endeavor is pointless. If there are indeed no hidden variables for predicting the outcomes of quantum measurements, then this is a much stronger argument than entropy for the arrow of time.

John Lindal
[BS '94, MS '95, PhD '01]

ANDREW LANGE REMEMBERED

I was saddened to hear of Professor Andrew Lange's passing recently. I was also disappointed to see that his obituary in the Winter 2010 *E&S* omitted his work with students of Caltech, and undergraduates in particular. More than 10 years on, I recall Lange's interest and involvement with the student body, through mentoring, student government and the House system. I'm sure that he would have wanted this mentioned, even in this tightly abbreviated listing of his contributions to the Caltech community. One online collection of folks remembering Lange through testimonials and photos can be found on Facebook; search for "Andrew Lange."

Kohl S. Gill [BS '98]

Professor Lange's obit, short by necessity, was adapted from the press release. A campus memorial service for him will be held on Friday, May 7. Look for complete coverage of it in the next issue of E&S. —ed.

ANNENBERG CENTER DEDICATED

I was pleased to read in the article about the dedication of the Annenberg Center that Caltech has a professor of computer science and applied mathematics.

In about 1976, I applied to study the application of then-nascent principles of computer science to the development of mathematical software. Professor Francis Buffington, who I had known as an undergraduate student lab technician in Engineering 91, had the duty to inform me that the computer science department (then called information science) thought my interests and qualifications fit well with the applied mathematics department, while the applied mathematics department thought I fit well with the computer science department.

Van Snyder, La Cresenta, CA E&S

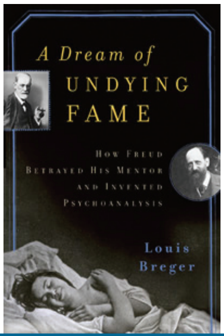

A Dream of Undying Fame: How Freud Betrayed His Mentor and Invented Psychoanalysis
by Louis Breger
Basic Books, 2009
160 pages, $22.95

BOOKS

This small biography is a gem, a brilliantly polished examination of a vital aspect of Sigmund Freud's life and work: his creation of the theory of psychoanalysis and, in the service of this accomplishment, his remarkable perfidy to his teacher, Josef Breuer, the physician who had mentored Freud and, notably, had introduced him to the "talking cure." Louis Breger, a psychoanalytic scholar, clinician, and Caltech professor of psychology emeritus, has published two previous biographies, *Freud: Darkness in the Midst of Vision*, which was a deep exploration of Freud's life and work, and *Dostoevsky: The Author as Psychoanalyst*.

As he did in his earlier Freud biography, Breger provides a nuanced, balanced, and ultimately respectful examination of the history of Freud's initial forays into psychotherapeutic treatment under the guidance of Breuer and the nascent psychoanalytic ideas they developed in collaboration. Breger's criticism of Freud's ideas is tempered by compassion for Freud's evident long-standing personal torment. Breger describes Freud's lifelong obsession with becoming famous and, to that end, his betrayal of Breuer, with whom he broke professionally and socially and whom he eventually failed to credit for their critically important early collaboration.

Breger's description of Freud's early history increases our understanding of Freud's development of some later, unsubstantiated theoretical views. For instance, Breger characterizes the conflict between Freud and his fiancée's mother and older brother:

"These fights were unconscious remnants of his reactions to all those sisters who took his mother from him as a child. . . ."

Confirmatory bias occurs throughout Freud's writings, and Breger brings it compassionately to life. He makes it clear that Freud's mind, brilliant as it could be in generating ideas, was simply not open to the possibility of disconfirmation of these same ideas. Although he initially framed his formulations as hypotheses, with the passage of time—but with no supporting evidence—he came to express these formulations as fact. He had to be right, and he required unwavering loyalty from his colleagues. This was a central tragedy limiting his greatness and his humanity.

Freud and Breuer claimed that cathartic recall of traumatic memories "immediately and permanently" relieved hysterical symptoms. However, as Breger notes, "it was much more difficult to achieve cures in practice than this statement implies." Contemporary evidence suggests that traumatic symptoms are less likely to be relieved by such clinical approaches than by methods intended to produce competing, benign memories. If Freud and Breuer were correct, one would expect, for instance, that the vivid, cathartic flashbacks of soldiers would be self-limiting, if not altogether curative, which they are not.

Since Freud is otherwise so vulnerable to confirmatory bias and other logical errors in the service of promoting his assertion that, for example, sexual factors underlie all forms of neurosis, one wonders how reliable Freud's claims might be of any of the details unearthed from patients' reports. Breger does not raise this specific concern, but, in the absence of confirmatory evidence, we can never know the accuracy of assertions such as Freud's claim that a young girl had been "sexually molested each night by her governess."

In fact, one might interpret the following to be a reflection of Freud's awareness of a fictional element to his case studies:

"It still strikes me myself as strange that the case histories I write should read like short stories and that, as one might say, they lack the serious stamp of science."

Freud's thinking was highly metaphorical—a fact that enlivened his writings but that also limited their basis in science. More, this dissociated his ideas from reality. He was, apparently, sometimes aware of this. Breger quotes a letter to Fliess, in which Freud confesses, "I no longer believe in my neurotica," adding that "there are no indications of reality in the unconscious." A contemporary critic of Freudian theory would be hard-pressed to express it more concisely.

Although Breger observes that it is this "fictional" quality to Freud's narrative style that enriches his writing and renders it so compelling, the question of its factual accuracy remains open. One glimpses the apparent ease with which a thought Freud may have had about his own psychology becomes a "fact" about everyone's. "A single idea of general value dawned on me. I have found, in my own case too, the phenomena of being in love with my mother and jealous of my father, and *I now consider it* a universal event in early childhood (italics mine)."

Thus, Freud admits that he takes an experience of childhood feelings of his own (taken by him as real, though we cannot know why) and then, as if by magic, he *considers* the experience to be a universal one. One can make the case that Freud's theories—his assertions about others' psychology—were essentially projections of his own sense of self. The fact that there is a kernel of truth to some of them reflects that his own psychology was not so different than that of most humans.

Breger's clinical acumen is beautifully rendered in his analysis of Oedipal theory. "Freud's substitution of his universal Oedipal theory for one based on real traumas was a mixture of truth and speculation. It revealed his wish for his mother's love and her loss to a rival, though he made the need for mother-infant attachment 'sexual' and substituted his father for the many babies who took his place. At the same time, it made him into a warrior, a young Oedipus, in combat with a king. It also did away with real traumas, sexual or any other kind, and gave primary emphasis to instincts and fantasies. In this new theory, it was not what actually happened that was the source of fear, depression, and symptoms—'the worries that robbed me of my youth'—but rather the young child's drive for pleasure, Oedipal fantasies, and sexual wishes that conflicted with moral standards. In addition, the theory itself—immediately promoted to 'universal' status—became Freud's bid for 'eternal fame'; it would make him a great scientist."

Breger's depiction of Freud's personality, especially his desperate quest for fame, shows us a pitiable man. For all of Freud's substantial talent at observation and, more, for his distinct and articulate expression, his personal torment was a profoundly limiting quality. With the benefit of contemporary psychological knowledge, one can readily see that many of Freud's ideas were based primarily upon his rich imagination. His talent for fantasy was surely as capacious as that of his hysterical patients. This commonality likely contributed to his extensive range of ideas, but not to their critical examination.

Breger suggests that Freud's *Studies in Hysteria* "began a revolution in our understanding of human personality and psychological disturbance." He adds that Freud "pushed the field in a number of fruitful directions." Freud's emphasis on both the existence and the ubiquity of unconscious motivation and on the potential meaningfulness of dreams continue to be powerful cultural influences. However, when Freud chose to not test his theoretical formulations with empirical research, he also led the field astray. Moreover, one might argue, he contributed significantly to the delay in the growth of the science of psychology, and especially to the delay in our understanding of personality and psychopathology. His overwhelming need for fame—and his corollary need, to be always correct—were, perhaps, satisfied. But his distortion of psychoanalysis from testable theory to a "cause" cast it into the category of dubious beliefs. "How different things would have been if, instead of a cult-like 'cause,' psychoanalysis had really been the science that it claimed to be. . . ." If only, Breger suggests, Freud could have continued to collaborate with Breuer.

This book is a remarkably successful depiction of a central aspect of Freud's life. Breger has written it as a scholar, yet it reads like a mystery, the solution to which is both compelling and tragic. *—JB* e&s

Joseph Barber is a clinical professor of rehabilitation medicine at the University of Washington School of Medicine in Seattle.

FINANCIAL SERVICES FOR THE GREATER GOOD®

Engineering & Science

California Institute of Technology
Pasadena, California 91125