LOS ALAMOS SCIENCE AND TECHNOLOGY MAGAZINE | JANUARY 2020

1663

Renewable nuclear energy
Asteroid interceptor
Algae-sucking vampires
Identifying infections

The Mind in the Machine

Los Alamos
NATIONAL LABORATORY
EST. 1943

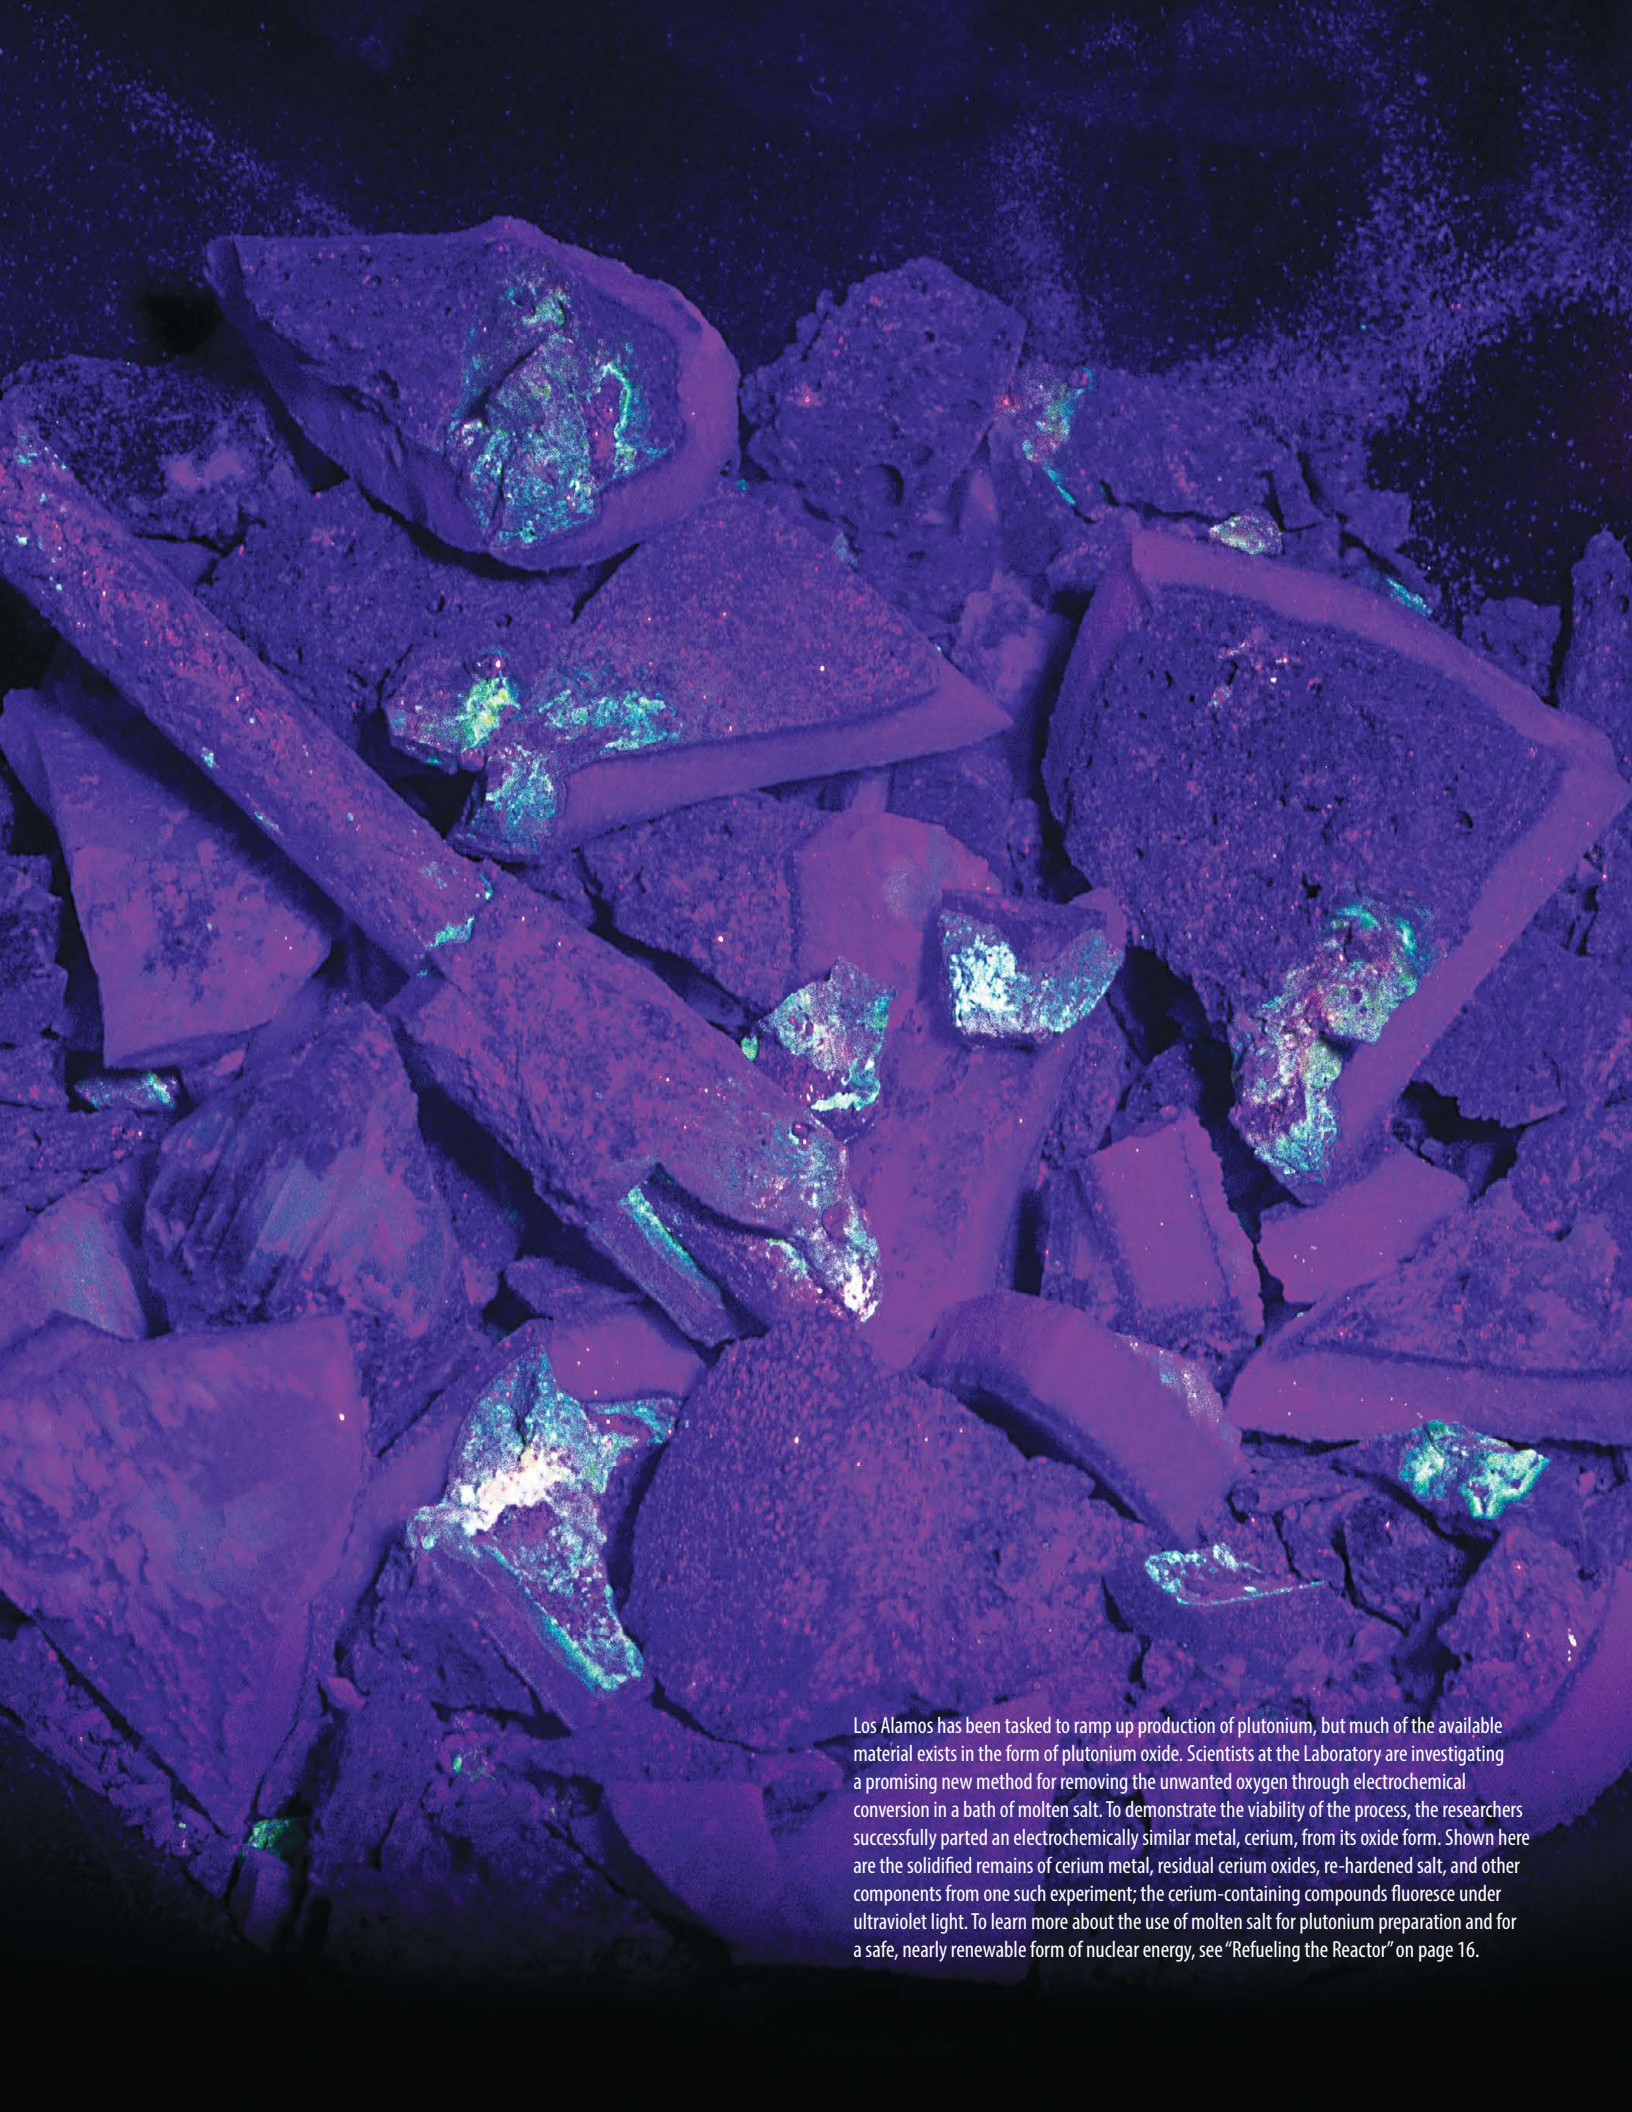

Los Alamos has been tasked to ramp up production of plutonium, but much of the available material exists in the form of plutonium oxide. Scientists at the Laboratory are investigating a promising new method for removing the unwanted oxygen through electrochemical conversion in a bath of molten salt. To demonstrate the viability of the process, the researchers successfully parted an electrochemically similar metal, cerium, from its oxide form. Shown here are the solidified remains of cerium metal, residual cerium oxides, re-hardened salt, and other components from one such experiment; the cerium-containing compounds fluoresce under ultraviolet light. To learn more about the use of molten salt for plutonium preparation and for a safe, nearly renewable form of nuclear energy, see "Refueling the Reactor" on page 16.

1663

LOS ALAMOS SCIENCE AND TECHNOLOGY MAGAZINE

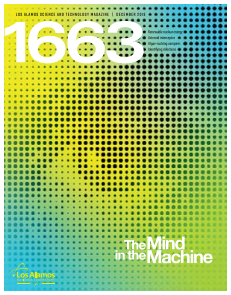

ABOUT THE COVER

Minds and machines have different skill sets. For example, minds do vision better than machines, but machines do computation better than minds. What can our minds do that we can teach machines to do? What can we learn from machines about our own minds? Machine learning—not quite artificial intelligence but much more than mere programming—is a rapidly advancing field now being explored by tech entities from social media companies to national laboratories. Scientists in all corners of Los Alamos are using and developing machine learning methods for a broad variety of national security challenges. Cataloging them all would be daunting, but what's key is that machine learning is revolutionizing how science is done, making some things that used to be impossible, possible.

ABOUT OUR NAME

During World War II, all that the outside world knew of Los Alamos and its top-secret laboratory was the mailing address—P. O. Box 1663, Santa Fe, New Mexico. That box number, still part of our address, symbolizes our historic role in the nation's service.

ABOUT THE LDRD LOGO

Laboratory Directed Research and Development (LDRD) is a competitive internal program by which Los Alamos National Laboratory is authorized by Congress to invest in research and development that is both highly innovative and vital to national interests. Whenever *1663* reports on research that received support from LDRD, this logo appears at the end of the article.

1663 STAFF

Craig Tyler Editor-in-Chief
Eleanor Hutterer Science Editor
Rebecca McDonald Science Writer
Dan Judge Designer
Sarah Tasseff Designer
Jathan Campbell Photographer
Daniel Owen Photographer
Kevin Sutton Photographer
Sheila Molony Copyeditor

2 SPOTLIGHTS

· Galactic positrons
· Lightning superbolts

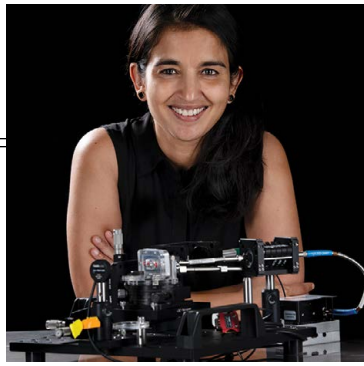

4 IN THEIR OWN WORDS

Harshini Mukundan is reinventing the detection of infectious diseases

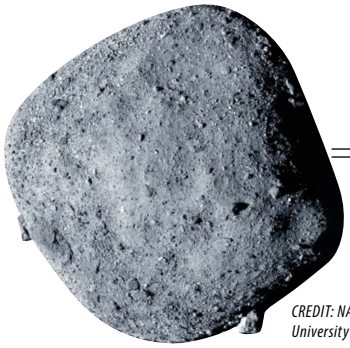

CREDIT: NASA/Goddard, University of Arizona

10 HOW TO SAVE THE WORLD

Two ways to prevent a catastrophic asteroid impact

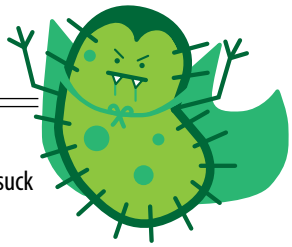

14 VANQUISHING VAMPIRES

Genomics to save algae ponds from bacteria that suck

16 REFUELING THE REACTOR

Safe, sustainable nuclear power with molten salts

24 THE MIND IN THE MACHINE

Machine learning for high-stakes science

A pair of nearby pulsars—ultradense, highly magnetized, and rapidly rotating neutron stars, shown here within a high-energy gamma-ray sky map—seemed a likely source for the excess of positrons flying through space and raining down on the earth. Recent observations, however, cast doubt on that hypothesis. *CREDIT: HAWC collaboration*

ASTROPHYSICS

Positively Perplexing

Why are there so many positrons in our galactic neighborhood?

Ten years ago, scientists observed something unexpected: an excess of positrons among cosmic rays. Where they're coming from is a galactic mystery that endures to this day.

Cosmic rays are energetic particles, mostly protons and atomic nuclei, that zip about the galaxy and rain down on Earth. Electrons and positrons, matter-antimatter partners having the same mass but equal and opposite electrical charges, are less common. But positrons in particular are substantially over-represented relative to theoretical models of their production in interstellar space.

What is producing them? Whatever it is, it must be within about a thousand light years of Earth, about 1 percent of the diameter of the galaxy. Otherwise, they would have lost too much energy from interactions with interstellar radiation and magnetic fields to account for their measured abundance at energies of tens to hundreds of gigaelectron-volts.

Theorists have posited only a handful of possible sources. Perhaps the most likely, or at least the most straightforward to investigate, would be a collection of astrophysical objects that produce and accelerate positrons. This would include pulsars, supernovae, and micro-quasars; the most likely of these, researchers expected, would be pulsars.

Pulsars are the ultradense relics of massive stars that went supernova sometime in the past. They have extremely powerful and rapidly rotating magnetic fields capable of acting as particle accelerators, and there are good reasons to expect that many of the particles will be positrons. Such accelerated positrons would collide with the photons of microwave light that permeate the universe, causing the microwaves to increase in energy and become teraelectron-volt (TeV) gamma rays. By observing the gamma rays coming from around a pulsar, one can see whether or not the pulsar is a significant contributor to the positron excess.

"We measured the TeV gamma rays from two nearby pulsars, both old enough for the cosmic rays to have reached Earth," says Los Alamos scientist Brenda Dingus. The results come from more than 500 days of observation using the High-Altitude Water Cherenkov Observatory (HAWC), a world-class TeV cosmic- and gamma-ray observatory operated by a Los Alamos-led international team. "We found the TeV emissions and confirmed that positrons are indeed being produced there," says Dingus. "And the emissions extended out several degrees from each pulsar, allowing us to measure how fast the positrons move away from the pulsar." Assuming that movement is indicative of how the positrons would travel to Earth, the HAWC team calculated the spectrum of positrons expected to be observable here.

The calculations revealed that the lower-energy positrons don't make it very far due to scattering by turbulence in the surrounding magnetic field, and the higher-energy positrons do not live long enough due to energy losses when interacting with the magnetic field. As a result, only positrons at a sweet-spot energy around 1 TeV could reach Earth—and not very many of them. But the observed positrons are numerous and span a range of energies lower than a TeV. It's just not a match. Unless positron propagation between the source and the earth is markedly different than anticipated for some unknown reason, pulsars are not the source. Or if they are, then they must be both nearby and, as yet, undiscovered.

The trouble is, if pulsars are wrong, then what is right? Alternative astrophysical objects—supernovae and micro-quasars—are still possible, even though the former are brief and the latter are rare. And it is always possible that there could be some completely unknown process of cosmic-ray propagation that generates the positrons. The only other significant possibility, according to theorists, would be dark matter as the source. If dark matter is made of a swarm of effectively invisible particles, then it's possible that the particles are susceptible either to decay, like a radioactive nucleus, or to collisions with one another. Either process would likely produce positrons. In that case, the source is effectively all around. Nothing in the HAWC data prohibits this.

The HAWC study helps chip away at an astrophysical mystery, but for the time being, the positron-excess anomaly and the search for a source—whether somewhere or everywhere—lives on. LDRD

—Craig Tyler

ATMOSPHERIC SCIENCE

Bigger, Badder Bolts

Lightning reveals its secrets only to the most painstaking analyses.

One-point-twenty-one gigawatts of 1980s film fame notwithstanding, the firmly established science of lightning is shockingly sparse. While the electrical arc itself is basically understood, the spark necessary to initiate it remains an area of speculation and active research (see "Out of Thin Air" in the March 2018 issue of *1663*). Just as puzzling is the fact that some intra-cloud bolts of lightning appear to be hundreds of times larger and more powerful than all the rest.

Los Alamos scientist Michael Peterson seeks out these "superbolts" and other novel lightning events by digging into colossal Earth-observation datasets from the National Oceanic and Atmospheric Administration's GOES-16 and GOES-17 weather satellites and the joint Los Alamos–Sandia national laboratories' FORTE nuclear-detonation detection and lightning-observing satellite. But unlike initial analyses of these satellites' data, which were optimized for rapid processing in order to provide real-time hazard warnings (GOES-16 and -17) or for recording single quick events (FORTE), Peterson developed a new algorithm to assemble a more complete picture of lightning in the historical data. In data from 2018 alone, his algorithm caught more than 14 million lightning events that had been underrepresented by the real-time data processing and were therefore widely ignored by the scientific community.

What Peterson found was downright astonishing. He discovered a population of flashes that streak horizontally over hundreds of kilometers. The longest of these was 673 kilometers, roughly the width of Kansas, while another spiderweb-like flash covered an area the size of Ohio.

Importantly, Peterson's observations of extreme lightning offer new insights into an old controversy, set off by a 1977 analysis of data from the Los Alamos Vela nuclear-detonation detection satellites. For decades, atmospheric scientists have argued about whether superbolts—defined as having 100 or more times the power of "normal" gigawatt lightning—are unique phenomena with their own distinct physics or simply the upper end of the distribution of normal lightning observations. The latter is a real possibility because the observed flash intensity depends greatly on the conditions under which it is viewed. For example, some clouds reflect additional light back to the camera, brightening the observed flash, while other dense clouds can get in the way and obscure the flash.

Peterson analyzed his data and found two scenarios that accounted for most of the observed superbolts. One of these involved the thunderstorm's "anvil" clouds. Around the edge of the storm, he reasoned, anvil clouds would present favorable viewing conditions, with reflecting layers that redirect light to the satellite. This would explain earlier findings that ordinary lightning with relatively weak electrical currents can still produce superbolt-class emissions and that the FORTE superbolts were less obscured by clouds than usual.

But the other population of superbolts was buried inside the expansive "stratiform" rainclouds that form adjacent to severe thunderstorms and are responsible for the long periods of light rainfall that linger after the storm passes. Stratiform clouds are arranged in uniform layers that enable unique lightning physics, and the most extreme cases Peterson identified were all examples of stratiform lightning.

So which is it? Unique hundred-gigawatt physics or just favorable viewing conditions?

"The simplest explanation would be one or the other," Peterson says. "But looking at the data, I'd have to say that it appears to be both. Some anvil superbolts aren't all that super, but the stratiform superbolts certainly appear to be. Their exceptional brightness is just another way in which stratiform lightning is unique. We still have much to learn about the physics that allows these beasts to be so powerful."

One-*hundred*-twenty-one gigawatts of power, at least.

—Craig Tyler

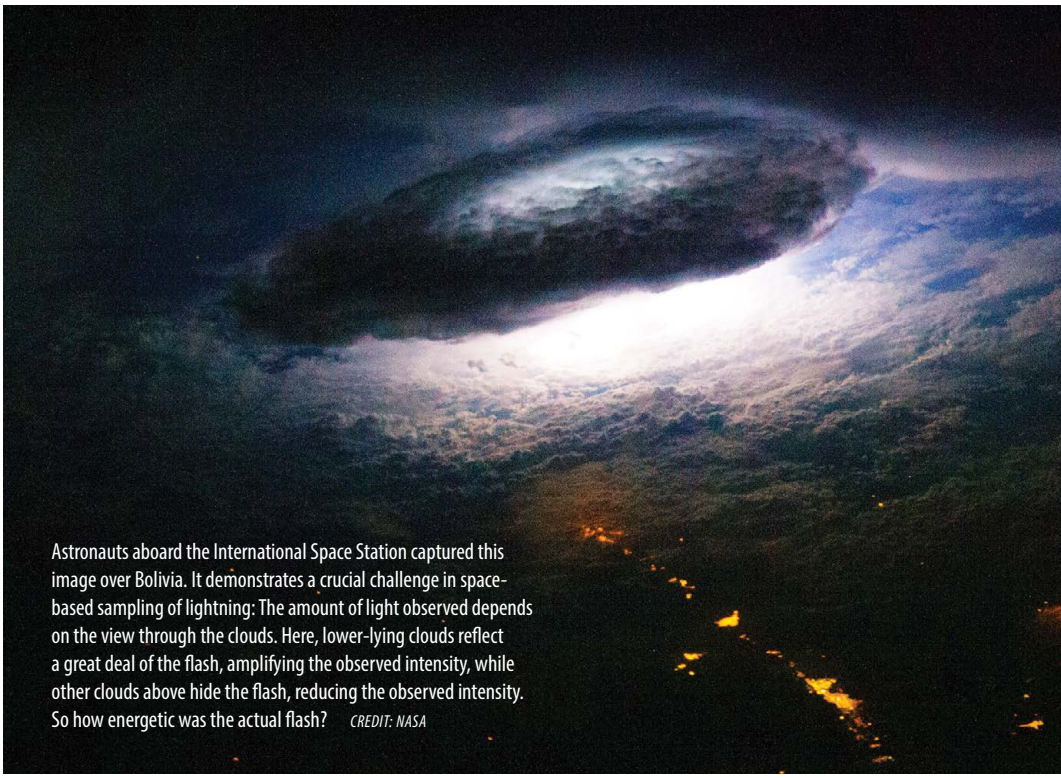

Astronauts aboard the International Space Station captured this image over Bolivia. It demonstrates a crucial challenge in space-based sampling of lightning: The amount of light observed depends on the view through the clouds. Here, lower-lying clouds reflect a great deal of the flash, amplifying the observed intensity, while other clouds above hide the flash, reducing the observed intensity. So how energetic was the actual flash? *CREDIT: NASA*

Thermo
HEPA

Mukundan (foreground) with four members of her team: (left to right) Zachary Stromberg, Loreen Stromberg, Kiersten Lenz, and Laura Lilley.

Reinventing Infection Detection

Microbiologist **Harshini Mukundan** relies on the chemistry of the immune system to diagnose disease quickly and accurately.

WHEN I VISITED KENYA IN 2016 as part of a new research collaboration, I took chocolate, candies, and toys to share with the children at the medical clinic and the local village. Growing up in rural areas, these children have limited exposure to the delicacies of the Western world—things that we sometimes take for granted—and bringing treats was a small, easy gesture that could put a smile on their faces. But my visit was part of a larger gesture—something my collaborators and I have been working on for many years—an effort to develop technologies to keep children healthy with increased access to medical advances.

Access to medical care is limited in resource-poor areas of the world, and families often travel great distances to seek treatment. Because of the economic burden, most families are unable to make multiple visits to the doctor. Hence, it is critical that the medical providers be able to diagnose and treat an infection right then and there at the point of need. This requires simple and effective diagnostic platforms that can give useful answers right away—answers that enable providers to respond quickly and to reliably dispense their limited supply of medications to the most needy patients.

First and foremost, medical professionals need to know: Is the causative agent bacterial or viral? And if possible, what specific pathogen is it? The challenge in addressing these questions is that there is no single test that can be used to diagnose all infections, as most diagnostic tools target only one specific type of pathogen. This technology gap forces clinicians to choose a diagnostic test based on details of the patient's symptoms—fever, congestion, vomiting—combined

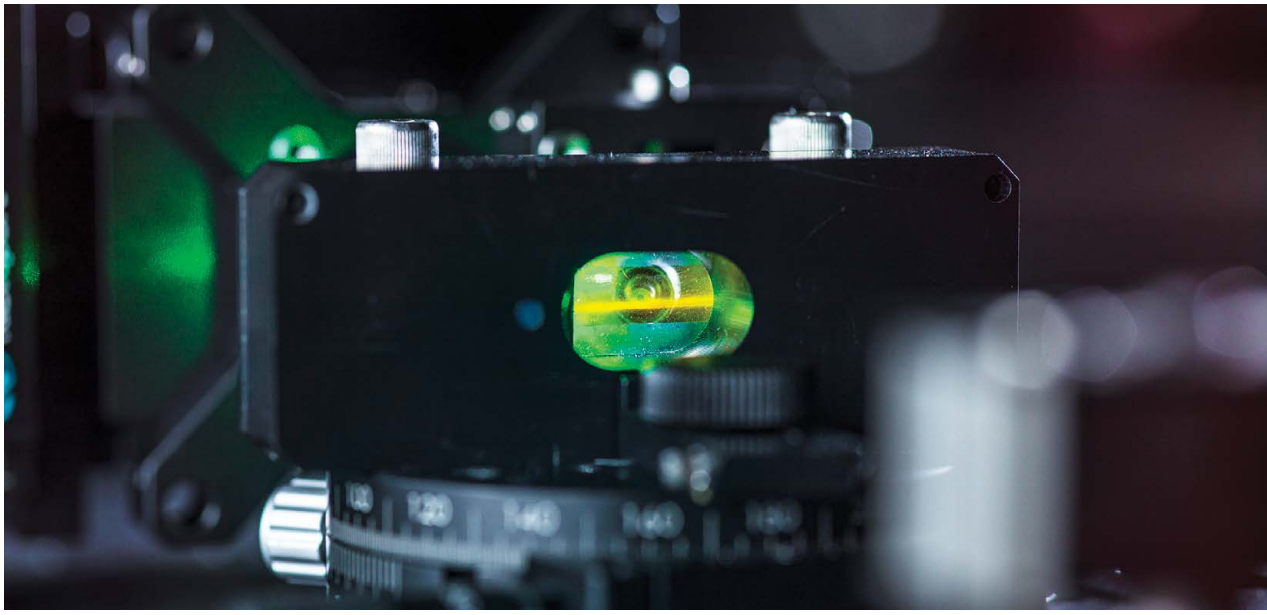

The optical waveguide sensor is a rapid, sensitive method for detecting certain molecules within a complex sample. Mukundan and her team have been adapting and miniaturizing the sensor platform for use in remote areas to detect bacterial pathogens.

with what diseases are currently prevalent in the area and what tests are available and affordable. Unfortunately, this so-called prior knowledge of the possible causes of disease inadvertently causes a bias in the diagnosis. Overcoming this situation and developing a universal diagnostic tool is my dream—and the challenge towards which our team has been working for the past decade. I am excited to say that we are making progress! We now have a prototype that can quickly screen samples for any type of pathogen—eliminating the need for prior knowledge—and that would be useful for doctors in first and third world countries as well as for veterinarians and even military personnel in rural areas who may not have access to extensive laboratory capabilities.

I get to tell this story, but the credit behind this progress goes to an incredible team of individuals at Los Alamos and collaborating institutions—a team of microbiologists, molecular biologists, immunologists, engineers, theorists, statisticians, chemists, and physicists who have worked with me toward this dream.

Immunity is innate

Growing up in India, I suffered from mumps and watched my sister battle measles. I was fascinated with infectious diseases and how the interaction between our immune systems and pathogens determines whether we suffer from a disease, defeat it, or succumb to it. As a high school student, my mandatory voluntary work—labeled "social useful productive work"—introduced me to tuberculosis patients at the local health center, where I learned that their treatment regimens had to be monitored for a grueling six months to ensure success. These events led me to think a lot about disease and medicine, which motivated me to study microbiology at university and pursue graduate studies in biomedical sciences in the United States.

I came to Los Alamos in 2006 as a postdoctoral researcher after finishing my Ph.D. in Biomedical Sciences at the University of New Mexico. My mentor, Basil Swanson, had been working on the development of a detection technique called a waveguide-based optical biosensor for the rapid detection of environmental organisms. Optical waveguides are small translucent plates made of two materials that differ in their ability to refract light. Propagation of light through the waveguide generates an optical field called an evanescent field, the intensity of which falls sharply as the distance from the waveguide increases. Thus, unlike in other platforms, the optical field is effectively confined to the surface of the waveguide where target molecules can be bound; this is an advantage because it eliminates the possibility of detecting extra "contaminant" molecules that may be present in a complex sample (other platforms would require additional steps to wash away these extra molecules). The optical waveguide confers excellent sensitivity and speed, although it does not add to the specificity of detection. However, when used in conjunction with fluorescently tagged molecules of interest (because the field is strong enough to excite the tags when they are bound to the waveguide surface), the combined detection technique is quick and effective.

I was awarded a National Institutes of Health postdoctoral research fellowship to explore the adaptation of this sensor technology toward the development of diagnostics for tuberculosis, one of the oldest and most challenging diseases known to man. Therein began my journey of trying to develop diagnostics. I was fortunate to have excellent mentors, collaborators, team members, and advisors—all of whom facilitated learning and advancement down this path.

During this time period, my infant nephew got sick with meningitis. Viral meningitis is self-limiting and patients usually

recover from it in about 10 days without treatment, whereas the bacterial form of the disease often requires extensive antibiotic treatment. To my surprise, I learned that there were no diagnostic tests to differentiate between the viral and bacterial disease. Thankfully, my nephew recovered and is now a strapping teenager. He did not require many rounds of antibiotics, so it is probable that he had the viral form and the antibiotics were not necessary. Yet the fact that he was treated with antibiotics simply because they could not differentiate the correct causative agent nagged at me, since unnecessary use of antibiotics can lead to antibiotic resistance. This led me toward a desire to develop a more universal diagnostic platform to discriminate bacterial infections from viral ones.

In our quest to develop such a platform, my team and I looked to the human immune system for inspiration. The human immune system has two parts: adaptive and innate. The adaptive system is well known for its development of antibodies that "remember" pathogens they've encountered before. On the other hand, the innate system is able to recognize and mount an immune response against invading pathogens effectively and quickly without any prior exposure to the specific pathogen. The innate immune system accomplishes this using a network of molecules that distinguish cells that belong to our bodies ("self") from foreign cells that don't belong ("non-self"). Some non-self molecules exist because they were released by the pathogen into the host during the course of infection. These molecules, known as biomarkers, are extremely consistent (or "conserved") across multiple pathogen species, and as a result, human immune receptors recognize them all as disease. Simply put, whether it is antibiotic-resistant tuberculosis (TB) or a newly emerging strain of staphylococcus, the innate immune system only needs to detect that one of these types of biomarkers is present to mount a response.

With this in mind, my team and I wanted to understand more about these conserved biomarkers and unravel the mechanisms by which they interact and associate with the human host. If we could mimic innate immune recognition in the laboratory, we could—in theory—repurpose it for a universal diagnostic strategy. Such a method would not require the user to have prior knowledge of what the pathogen could be, would not be driven by the patient's symptoms, could be applied for existing and emerging infections with equal efficacy, and could provide early diagnostic information to guide decision making and suitable therapeutic intervention.

We started by working on such a strategy for the diagnosis of bacterial infections, specifically TB.

LAM and the biological taxi service

Tuberculosis, caused by the bacterium *Mycobacterium tuberculosis*, is one of the oldest diseases known to infect humans, yet it is notoriously difficult to diagnose. With the evolution of drug resistance, and the recent phenomenon that tuberculosis is often associated with HIV infection, the diagnostics problem has become even more severe in the past few decades.

Several investigators have considered detecting one of these aforementioned biomarkers: a molecule called lipoarabinomannan (LAM) that appears in patients with an active TB infection. A component of the bacteria's cell membranes, LAM is a lipidated sugar, or lipoglycan, meaning that it has a sugary part and a fatty part. When *M. tuberculosis* cells are engulfed by some of our immune cells (macrophages) in the lung, the TB cells encounter innate immune receptors. LAM is known as a virulence factor because it activates the immune receptors in this environment, and therefore direct measurement of LAM in infected patients can provide an effective strategy for diagnosing active TB.

Several investigators (including us) have developed tests to measure LAM in urine, but uric acid can break apart the LAM molecule, so detecting it in blood would be better—albeit more difficult. Upon investigation, my team and I realized that the difficulty was directly owing to the biochemistry of LAM. LAM is an amphiphile, meaning the lipid part is hydrophobic (water-repelling) and the sugar part is hydrophilic (water-loving). Therefore, like oil droplets that group together in water, hydrophobic molecules find each other rather than float around freely in aqueous blood. However, because the human body is composed of many hydrophobic lipids that often need to travel in blood, our bodies have lipid-carrying molecules, called lipoproteins, whose job it is to transport lipids from one part of the body to another, behaving as a "biological taxi service." Examples of these courier molecules include high-density and low-density lipoproteins (HDL and LDL), which are commonly known as indicators of cardiac health. My team and I wondered if these lipid couriers also transported pathogen biomarkers in blood, and exploration of this hypothesis resulted in a resounding "yes"! HDL and LDL function as carriers for both host and pathogen lipids through the blood.

In our quest to develop a universal diagnostic platform, we looked to the human immune system for inspiration.

Upon further study we learned that although LAM is a biomarker for TB, most bacterial pathogens secrete other kinds of lipidated sugar biomarkers that are involved in virulence and immune recognition. For instance, one class of bacteria, called gram negative, releases lipopolysaccharides (LPS), and another class, gram positive, releases lipoteichoic acid (LTA). The host's HDL, LDL, and other lipoprotein molecules transport these biomarkers around.

Unraveling this concept of host-pathogen interactions helped solidify a strategy for my team: Using LAM, LPS, LTA, and other biomarkers, we developed two detection assays by capitalizing on their association with HDL and LDL carrier molecules. This is a relatively novel approach because many current diagnostics instead rely on detecting protein-based molecules that are hydrophilic and thus easily found in blood or urine.

Our first assay, lipoprotein capture, tested the concept of using lipoproteins but in a way that requires prior knowledge of the target pathogen, making this assay most useful for clinical research and animal studies. Based on the success of capturing and detecting lipoproteins, we developed our second assay, membrane insertion, to be a truly universal diagnostic approach:

one that does not require any knowledge whatsoever of the pathogen or its interaction with the host. Using the membrane insertion assay, a clinician would be able to quickly and accurately determine if any type of bacteria is present in the sample, and if not, they could safely assume that the infection must be caused by a virus or parasite. This assay capitalizes on the knowledge of the bacteria-specific biomarkers: LPS for gram negative, LTA for gram positive, and LAM for TB.

The first step of the membrane insertion assay requires blood-sample processing (which takes about two minutes) to separate the biomarkers from their HDL and LDL carriers. Second, we prepare a waveguide that has a lipid bilayer on its surface, mimicking a cell membrane. Once separated from their lipoprotein carriers, the hydrophobic ends of the biomarkers are attracted to the hydrophobic part of the lipid bilayer and thus insert themselves into the membrane. Next, fluorescently labeled antibodies that specifically target LPS, LTA, or LAM are introduced and—if they to bind to the biomarkers now held close to the waveguide—optically detected with the biosensor to confirm the presence and identity of the biomarkers. We have also demonstrated that this method can be multiplexed to measure many things at once by using different fluorescent tags on each of the various antibodies.

Furthermore, because the sample processing we developed is universally able to separate any kind of amphiphilic biomarker molecule from multiple kinds of host carriers, this approach allows for a one-size-fits-all strategy for the diagnosis of bacterial infection, bringing us a step closer towards achieving our goal of a universal bacterial sensor.

Fit for travel

Molecular assays and sensitive detection strategies are only of real value when successfully applied to a pressing clinical problem. From the very beginning, our team has collaborated with the National Institutes of Health, Johns Hopkins University, Medical University of South Carolina,

Universal Bacterial Sensor

Mukundan and her team have developed a new strategy for quickly and accurately detecting the presence of a bacterial infection.

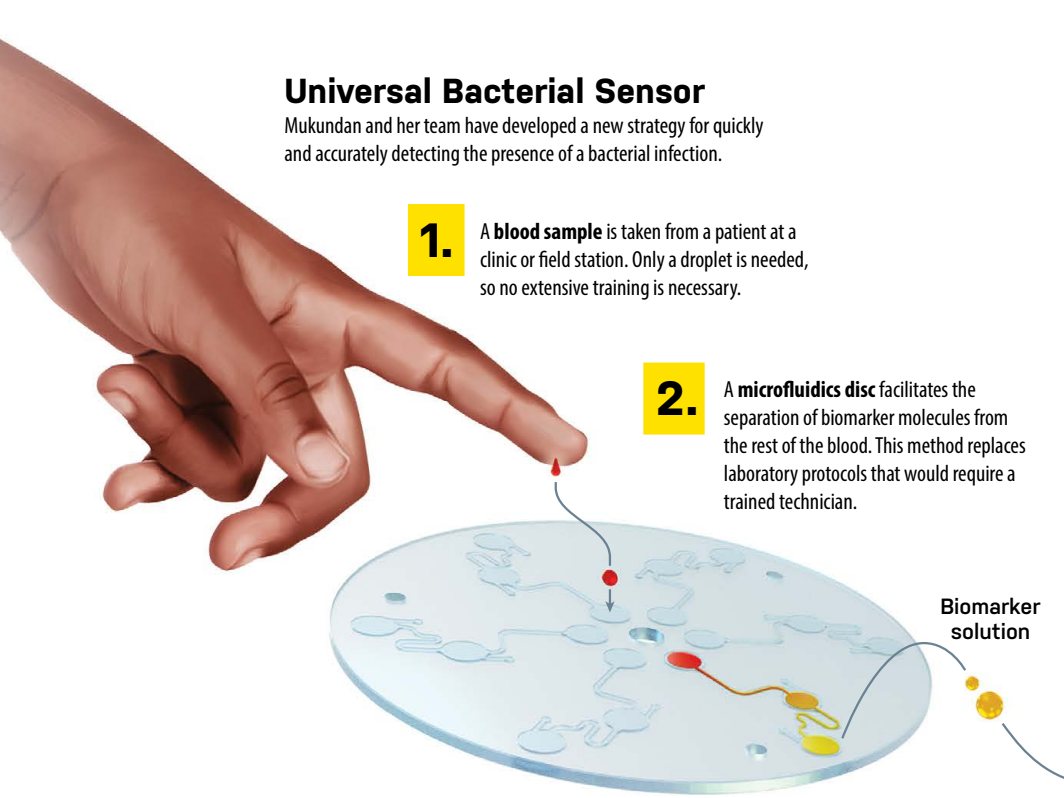

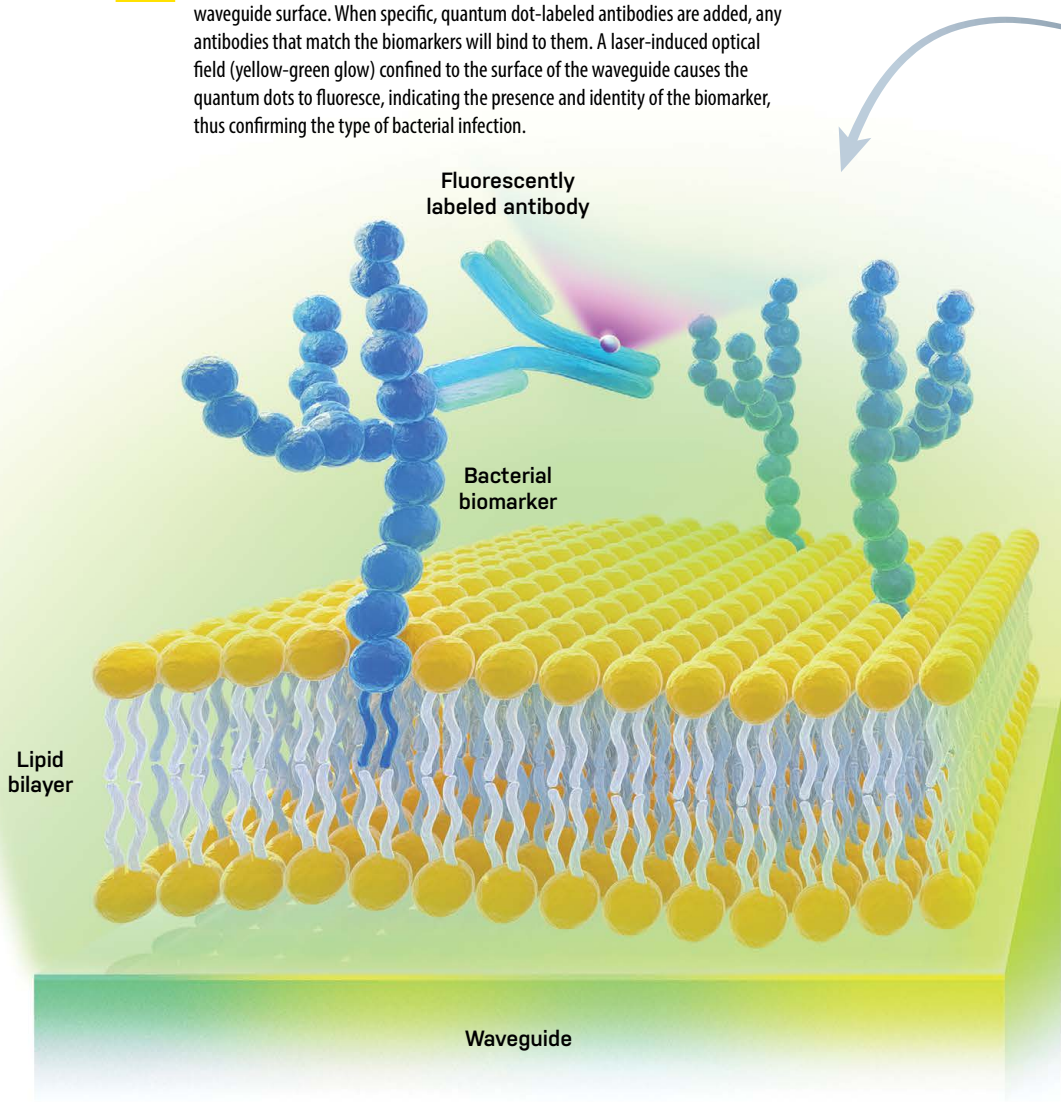

and other institutions to adapt our assays to the clinical scene. In 2015, theoretical biophysicist Ben McMahon and I, along with others in the Los Alamos Theoretical and Bioscience divisions, began a collaboration with Dr. D. J. Perkins and his team at the University of New Mexico Center for Global Health to evaluate our diagnostic methods with pediatric clinical samples from a resource-poor clinic in Siaya, Kenya. This collaboration has proved invaluable and fruitful in many ways.

In sub-Saharan Africa, a leading cause of childhood mortality is bacteremia, also known as sepsis: the condition when bacteria are present in the bloodstream. Unfortunately, this condition (bacteremia itself is not a disease) is often difficult to diagnose; although microbial culture is the gold standard, it is slow and not particularly sensitive. In addition, patients often have co-infections such as malaria or HIV that can complicate the diagnosis. Our universal sensor could help these clinics by enabling them to easily screen patients for bacterial infection—providing diagnostic information in a timely manner to guide treatment. However, a few things needed to be modified for our methodology to be applicable in this scenario.

3. The biomarkers are added to a **waveguide** that has been prepared with a lipid bilayer to mimic the membrane of a cell.

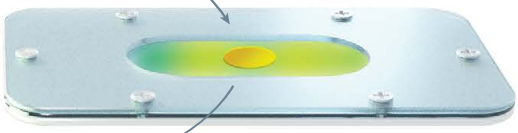

We want to make it easy for clinicians to determine if any type of bacteria is present in the sample.

For one, although our processing steps to separate the serum and biomarkers from the blood are quick, they require laboratory infrastructure not readily available in resource-poor areas such as rural Africa. To resolve this, we have been working on developing a microfluidics approach, which will enable rapid, lab-free sample processing at the point of need. Collaborating with experts in the Physics and Bioscience divisions at Los Alamos, we have designed and validated a microfluidics disk and are currently integrating it with the sensor platform. This is the first time a microfluidics approach has been used for lipid extraction from blood, and the outcome stands to be a simple and safe system that could facilitate the deployment of our sensing platform.

In addition to optimizing the sample preparation strategy, we also needed to re-evaluate the waveguide-based sensor platform. The waveguide platform has provided greater detection sensitivity compared to conventional methods in all our previous evaluations; however, it was originally developed over a decade ago and is not suitable for use in the field. Therefore, we decided to simplify and reengineer the platform to make it travel-friendly. In collaboration with scientists from the Physics Division and engineers at the Los Alamos National Security Education Center, we have now successfully miniaturized the instrument and developed a portable version; we are currently working on evaluating it and optimizing its sensitivity and performance.

The whole idea is to have one box that anyone can use at the point of need: a primary care doctor or specialist, a soldier, an emergency responder, or a veterinarian. And I'm happy to say we're getting close to having one.

Lessons in life and science

Working on this biosensor has taught me a number of things about chemistry and medicine and has exposed me to new challenges in optics, engineering, and informatics. But beyond the science, this effort has taught me a great deal about teamwork. Bringing the universal bacterial sensor to this point of development has required input from engineers, physicists, chemists, informatics experts and theorists, biologists and microbiologists, clinicians, molecular biologists, and veterinarians—a truly multi-disciplinary collaborative effort. The project has allowed each of us the opportunity to learn something new, to contribute something to the final product, and to be part of a wonderful team. The team has also included students, postdoctoral fellows, technologists, and engineers—each of whom has a unique perspective that helped shape the science. Not only did the team demonstrate intellectual diversity, but also social diversity—involving individuals from multiple social, ethnic, and gender backgrounds from various countries, including a team in Kenya that has played a critical role in the clinical recruitment and evaluation.

This article may be focused on my impressions—but these impressions and thoughts are made possible by the enthusiastic and passionate participation of members of this team. As a team, we hope to make an impact on diagnostics in the future. Our goal—indeed, our dream—is to remove any element of guesswork from the diagnosis of infectious disease and to have specific and rapid identification of a pathogen at the point of need. Realizing this dream is certainly a work in progress that has had many disappointments and derailments along the way. But there are two lessons I like to live by: One is to never give up on the things you really want in life, and the other is to not be afraid to change direction when something does not go the way you planned or envisioned.

So onward we go, without giving up, changing direction as needed along the way. One of the best aspects of this journey is that it is so easy to be inspired by the children in Kenya. These little ones are by far the most positive and cheerful individuals I have encountered in my life. On one of these visits, I would like to take something more than chocolates: the hope for a healthy future. LDRD

—Harshini Mukundan

HOW TO SAVE THE WORLD

An asteroid impact is the one natural disaster we can actually prevent.

THE NASA ADMINISTRATOR GETS THE CALL. A new asteroid has been detected. It's big—not drove-the-dinosaurs-to-extinction big, but still nearly a kilometer across. Its trajectory suggests that, in a few years, it could crash into the earth. If it does, it will completely devastate some part of the world: obliterate a major city, flatten a forest, bury a huge swath of land in fiery rock, or, if it hits at sea, potentially wipe out hundreds of miles of coastline with a tsunami.

There are thousands of asteroids with Earth-crossing orbits, and close calls are not so rare. Just this past July, an asteroid large enough to destroy a city passed within about 75,000 kilometers (km) of the earth; that's only about five earth-widths away and less than one fifth of the average distance to the moon. Had it crossed directly ahead of the earth's path, it would have missed by only 42 minutes.

There are a handful of people in the world making preparations to defend the world against killer asteroids. Among them, Cathy Plesko and her collaborators at Los Alamos—colleagues, postdocs, and students—are working out the plans necessary to intercept an incoming asteroid and nudge it off course with as little advance notice as possible. As things stand today, Plesko thinks we could develop and implement a plan to deflect a large asteroid if we had five or ten years of lead time—time to develop and launch a mission, time for the spacecraft to reach the asteroid, and time for the redirected asteroid to edge far enough off course to skirt around our planet.

"That far out, we couldn't gauge its orbital path with enough accuracy to know for certain that it will hit," says Plesko.

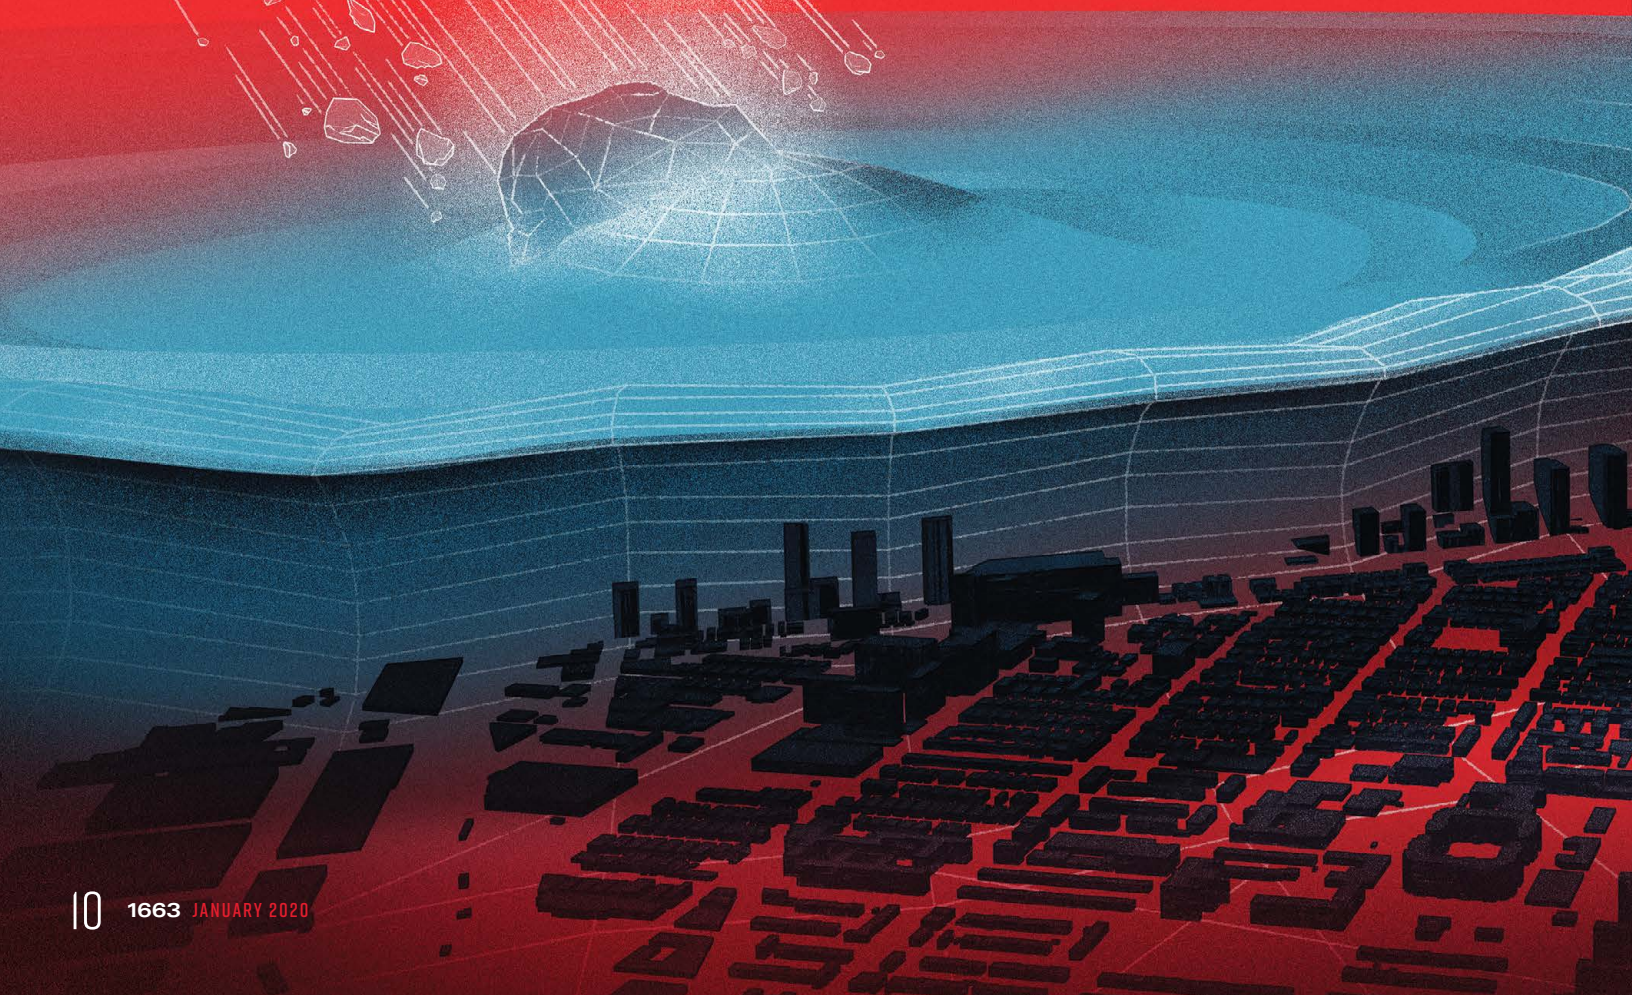

"There might only be a one-in-four chance. But if we wait long enough to become fairly certain, there won't be enough time to act. It's a difficult problem."

Chelyabinsk and Chicxulub

Different asteroids threaten different amounts of damage, depending largely on their size. In 2013, a 20-meter (20-m) diameter asteroid exploded about 30 km above Chelyabinsk, Russia, well above most of the atmosphere. The explosion was reportedly brighter than the sun and produced a shock wave that arrived on the ground several minutes later, breaking glass and causing other damage to thousands of buildings in the dead of winter. In the shock wave and the subsequent panic, about 1500 people were injured. Although no one was killed, the explosion produced about 30 times more energy than the Hiroshima atomic bomb.

The most powerful meteoric airburst ever recorded was a little over a century earlier—also, coincidentally, over Russia. That object was probably several times larger than the Chelyabinsk asteroid and its explosion perhaps 50 times more energetic (estimates vary). It is believed to have penetrated to less than 10 km above the ground, where its fireball flattened 2000 square kilometers of forest and killed several people.

Such airburst events, from asteroids in the tens or low hundreds of meters in diameter, would be difficult to prevent. By virtue of their small size, these objects reflect little sunlight and can therefore be virtually undetectable; the Chelyabinsk object, for example, was unknown prior to its arrival. Furthermore, the level of damage caused by such events is limited enough that it may not justify the expense of a space mission to prevent it.

At the other end of the spectrum, an incoming object several tens of kilometers in diameter, such as the one that produced the Chicxulub crater in the Mexican Yucatán and is believed to have caused the mass extinction that wiped out the dinosaurs, would be very difficult to deflect from a collision course because of its sheer inertia. Perhaps it could be done, but it would probably require a large number of space missions to do it.

Plesko has so far focused on the fertile ground in between: objects ranging from several hundred meters to several kilometers in diameter. Such objects are more numerous than the extinction-causing ones and generally survive the trip through the atmosphere and reach the ground. A best-case scenario would be an impact in the middle of the ocean, far from land. A series of large circular waves would expand outward in all directions, attenuating as they travel, perhaps generating a small tsunami on the closest shorelines. If the same event occurred within a few hundred kilometers of land, however, the effects would be devastating. Plesko's colleague Galen Gisler developed a computer simulation that showed such an impact would produce waves hundreds or thousands of meters high. (The tallest building in the Western hemisphere, One World Trade Center in New York, stands at 541 m.)

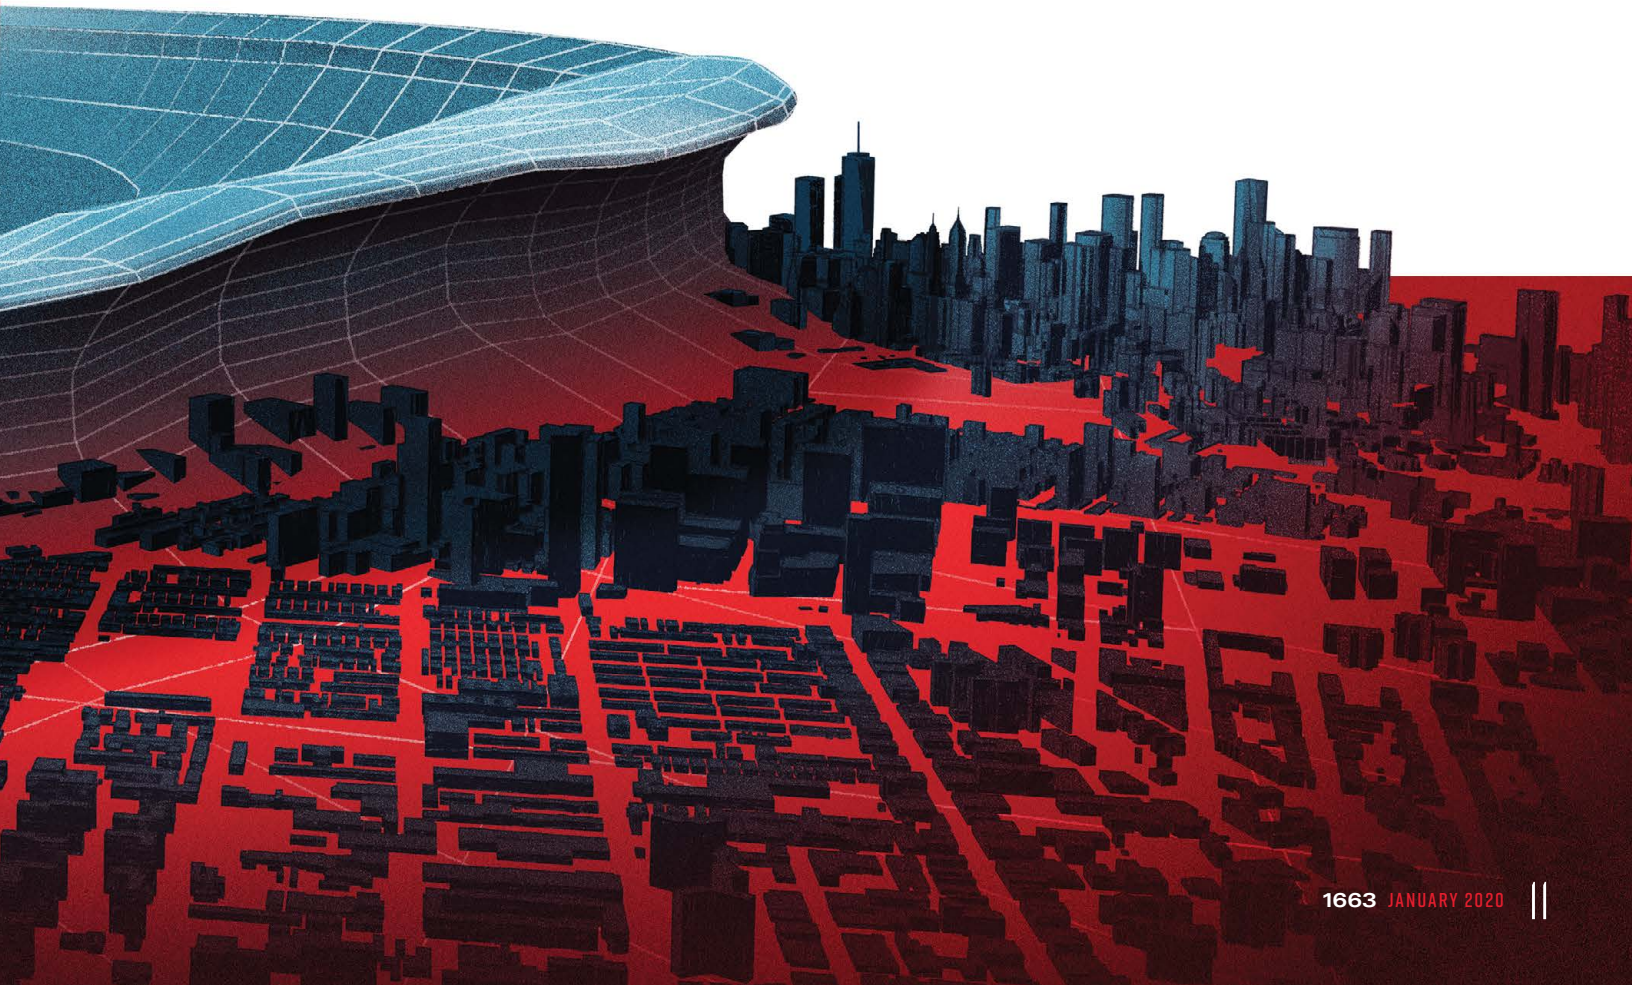

(Above) The asteroid Bennu is currently being orbited by NASA's OSIRIS-REx spacecraft, which will eventually touch down, collect a sample, and return the sample to Earth for study. Bennu could threaten a collision with the earth in the 22nd century; understanding its composition will help scientists determine how best to deflect it, and others like it, off of a collision course. (Right) NASA's DART mission will experiment with intercepting and deflecting an asteroid in the binary-asteroid system Didymos. A kinetic impactor (weighted spacecraft) will slam into "Didymoon," the smaller of the pair, and thereby alter its orbit around the larger body without changing the likelihood of either one approaching Earth in the future.

CREDIT: (Above) NASA/Goddard, University of Arizona, Lockheed Martin; (right) NASA

A similar object striking land would result in a crater ranging from a few kilometers in diameter to several tens of kilometers. For example, an asteroid a few kilometers across could produce a crater roughly 20–30 km wide—several times the size of Washington, D.C.—and eject a massive amount of rock, soil, and other debris, burying everything over hundreds of kilometers in every direction.

This is what one rocket, or maybe a few, might prevent.

Plan A: ram it

Plesko uses computer models to evaluate the effectiveness of potential space interventions to deflect asteroids off of a collision course. There are two options: a nonnuclear option and a nuclear option. The nonnuclear option is called a kinetic impactor; the spacecraft itself, heavily weighted, rams the asteroid. (The nuclear option, in this context, isn't metaphorical; it's an actual nuclear detonation.)

"On the face of it, the kinetic impactor is severely limited by how much weight we could get off the ground and then get up to speed," says Plesko. "But here the devil is in the details. Depending on the composition of the asteroid in question, we might get a serious enhancement effect on impact."

It may seem counterintuitive, but simulations show that the momentum ultimately imparted onto the asteroid can be significantly larger than the momentum of the impactor itself. When the impactor strikes, it causes a great deal of ejecta to blast off of the impact surface, producing an equal and opposite recoil. As a result, there is a gain relative to the kinetic impactor strike. How much of a gain depends on the "competence" of the asteroid. For a single, solid (competent) chunk of rock, there will be little ejecta; the gain might be 20 percent. For a loose assembly of smaller rocks held together only by the relatively weak gravity of the asteroid, there could be a great deal of ejecta, and the gain might be as much as ten times the original momentum of the impactor.

> If you hit a rock **hard enough,** it **flows like water.**

However, even if the asteroid is loosely bound together, simply breaking it apart with an impactor may not be ideal; the details matter. If the asteroid produces a modest amount of high-speed ejecta and the remaining body is deflected off course, then great. But if the asteroid essentially disintegrates but remains on course for Earth, then the resulting spray of smaller objects could still do a great deal of damage by shredding satellites and by heating and dust-loading the atmosphere, resulting in various complex and destructive climate processes.

Plesko works with Los Alamos mathematical physicist Len Margolin. Together, they build and run simulations on kinetic impactor outcomes. First the impact compresses the asteroid and vaporizes part of it. This produces an explosion and a pressure wave rippling through the asteroid body and resulting in ejecta launching from across a wide stretch of the asteroid's surface centered on the point of impact. To understand the transmission of forces through the asteroid, the simulation captures not only the physical properties of rock—an area that Plesko, a geophysicist by training, holds near and dear to her heart—but also the detailed physics of fluid flows.

"If you hit a rock hard enough, it flows like water," says Plesko. The simulation she and Margolin built treats the asteroid accordingly, breaking it up into a large number of tiny fluid cells, like compressible 3D pixels. But instead of the pixels having values for red, green, and blue, they have values for pressure, temperature, and other fluid properties, and a supercomputer tracks how forces are transmitted from one cell to the next. By virtue of its experience with nuclear-weapons simulations, Los Alamos has tremendous expertise in this kind of computer modeling.

Bennu and Didymoon

Computer simulations are only as good as the physical data fed into them, and here, Plesko and Margolin are getting help from NASA on the biggest unknown factor, the composition of the asteroid. In most incoming threat cases, this will be unknown; but it may be possible to either compare telescope observations

with data from other asteroid-visiting space missions in order to make an educated guess or, if there's enough lead time, launch an earlier spacecraft to study the asteroid before settling on the trajectory and other details of the intercept mission.

A NASA spacecraft called OSIRIS-REx is currently orbiting an 800-m asteroid called Bennu. Bennu passes near Earth every six years but is not expected to threaten a collision until sometime in the next century, possibly. However, Bennu is considered representative of a class of dangerous asteroids, and OSIRIS-REx is will collect some material from the asteroid in 2021 and fly it back to Earth for scientific study.

Meanwhile, another NASA mission will actually test a kinetic impactor on an asteroid. The Double Asteroid Redirection Test, or DART, for which Plesko is an active collaborator, will launch in 2021 and visit a binary asteroid system called Didymos. Within Didymos, a smaller, 160-m asteroid, affectionately but unofficially called "Didymoon," is bound in orbit with a larger, 780-m one. DART will converge with Didymoon in 2022 on a trajectory designed to alter Didymoon's orbit around the larger body without changing either body's orbit around the sun. Data collected by ground-based observations will be used to evaluate the effectiveness of kinetic impactors.

Both OSIRIS-REx and DART will return valuable information to help constrain the major unknowns in the Los Alamos simulations and calibrate expectations with hard data. Undoubtedly, that will sharpen the line between those incoming asteroids that can be effectively handled by kinetic impactors and those that cannot.

Plan B: fry it

The ideal course of action to deflect an incoming asteroid depends on many factors, such as its size, competence, and orbital trajectory—and how much time remains before it hits. For the right kind of asteroid, kinetic impactors are appealing because of their simplicity: a large mass attached to a rocket. But if the asteroid is too large or there isn't enough time, and the only way to save the world is by delivering a lot of energy to the asteroid as quickly as possible (rather than launching a series of kinetic impactors, say), then a nuclear explosion is the only way to do it.

To assess the nuclear option, Plesko collaborates with the Lab's Steve Becker. Their simulations have demonstrated two promising approaches. The first is the obvious one: fly right up to the asteroid and detonate the weapon on it. This "disrupt and disperse" approach is suitable when there isn't time for anything else and, as Plesko puts it, "you just have to get rid of the sucker."

But a more promising nuclear option, the simulations reveal, would be a nuclear detonation near, but not actually on, the asteroid. The explosion would produce a blast of energetic x-rays, which would immediately vaporize, or ablate, the surface of the asteroid. The resulting expanding gas would produce a powerful recoil, driving the asteroid away without creating a lot of dangerous debris. How far away to detonate depends on two competing factors; the closer the detonation, the more energy is directed at the asteroid rather than empty space, but the farther away, the more of the asteroid's surface will be exposed to x-rays. The ideal distance strikes a compromise between these two effects, and Plesko and her colleagues can calculate approximately how far from the surface that "sweet spot" lies.

An additional benefit of this ablation-from-a-distance method is that it spreads out the pressure on the asteroid, pushing evenly across a wide surface (like a shove), rather than concentrating all the force on one spot (like a stab), as a surface detonation or a kinetic impactor would do. In fact, even for incoming asteroids with size and lead time suitable for a kinetic impactor, nuclear ablation may still be the way to go if the competence of the asteroid is in question, as it often is.

Be prepared

Plesko has the simulation producing realistic results. Two NASA missions will provide important calibrating data. It then remains to examine the simulation under a variety of conditions: various incoming trajectories, shapes, sizes, masses, and compositions. So far, she has focused on roughly kilometer-scale asteroids; she will need to broaden that focus to include larger objects (like the one that caused Chicxulub) and comets (which are not made of rock). The goal is to have a set of ready responses for different classes of incoming objects. It would also help to build one or more rockets in advance. If the hardware is already in place and allows a reasonable degree of operational flexibility, then humanity can shave years off the necessary lead time: spot a threat, run the simulation, identify an intercept trajectory, load either a warhead or a kinetic impactor mass, and start the countdown.

The **blast of ejecta** produces an **equal and opposite** recoil in the **asteroid body.**

Humanity lives now much as it ever has, at the mercy of numerous types of natural disasters. Tornadoes. Hurricanes. Volcanoes. Earthquakes. Yet unexpectedly, a catastrophic meteor strike is the one that's technologically preventable, given enough preparation. And rapidly spooling up new technology to address an urgent threat—well, that's a big part of what Los Alamos is known for. LDRD

—Craig Tyler

MORE **PLANETARY DEFENSE** AT LOS ALAMOS

http://www.lanl.gov/discover/publications/1663/archive.php

- **Protecting satellites from space junk**
  "Averting Orbital Apocalypse" | July 2019
- **Collision-course comets**
  "Diverting Doomsday" | December 2016
- **Stormy space weather**
  "The Stuff that DREAM Is Made Of" | June 2012

from **1663**

Vanquishing Vampires

Genomic analysis of predatory bacteria stands to improve the future of biofuels.

Floating together in a vast artificial pond, millions of vibrant green algae represent a bright future—one with an abundance of sustainable vehicles fueled by the oil these tiny organisms produce. Lurking below, however, vampire-like bacteria threaten to destroy the algae by sucking out their juicy insides. *Vampirovibrio chlorellavorus* are predatory bacteria, discovered in 1972, that specifically target algae in the *Chlorella* genus, which happen to be some of the best-suited algae for growth as a biofuel feedstock.

Although *V. chlorellavorus* can be eliminated using biocide chemicals or acids in order to save the algae, scientists at Los Alamos and the University of Arizona are studying their bacterial genomes to better understand how these and other pests may threaten algae ponds in the future. Recent analyses suggest that there may be multiple species of *Vampirovibrio* that prey on *Chlorella*. Furthermore, it is unknown if these new species or strains respond to the known treatments, and if they don't, well, for biofuel scientists, that might be like finding vampires who are not repelled by garlic.

"We're just scratching the surface here," explains biologist Blake Hovde. "As we move algae production outdoors to be more economically feasible, there will be additional unknown pests out there. We need a standard way to identify them so they can be removed quickly."

Hovde and his team at Los Alamos have been evaluating the genomes of these invasive bacteria by isolating them from pond samples collected by collaborators at the University of Arizona. The analysis includes identifying organisms that harbor known virulence and pathogenicity genes, but it also includes discovering new genes that cause disease in algae, which may help the researchers identify new organisms of concern. They have also been studying a bacterial pathogen that targets *Nannochloropsis*, which is another genus of algae favored for biofuel production.

Identification of these specific gene targets allows the Los Alamos team to develop primers, which are small sections of DNA that can be used to target that same gene in a new sample. Ultimately, the Los Alamos team is developing a standardized way for algae farmers to test their ponds on a daily basis using field-ready equipment coupled with a set of primers that target all known predatory organisms. Rapid, specific detection of the exact threat will enable farmers to respond with the appropriate treatment quickly before it's too late.

—Rebecca McDonald

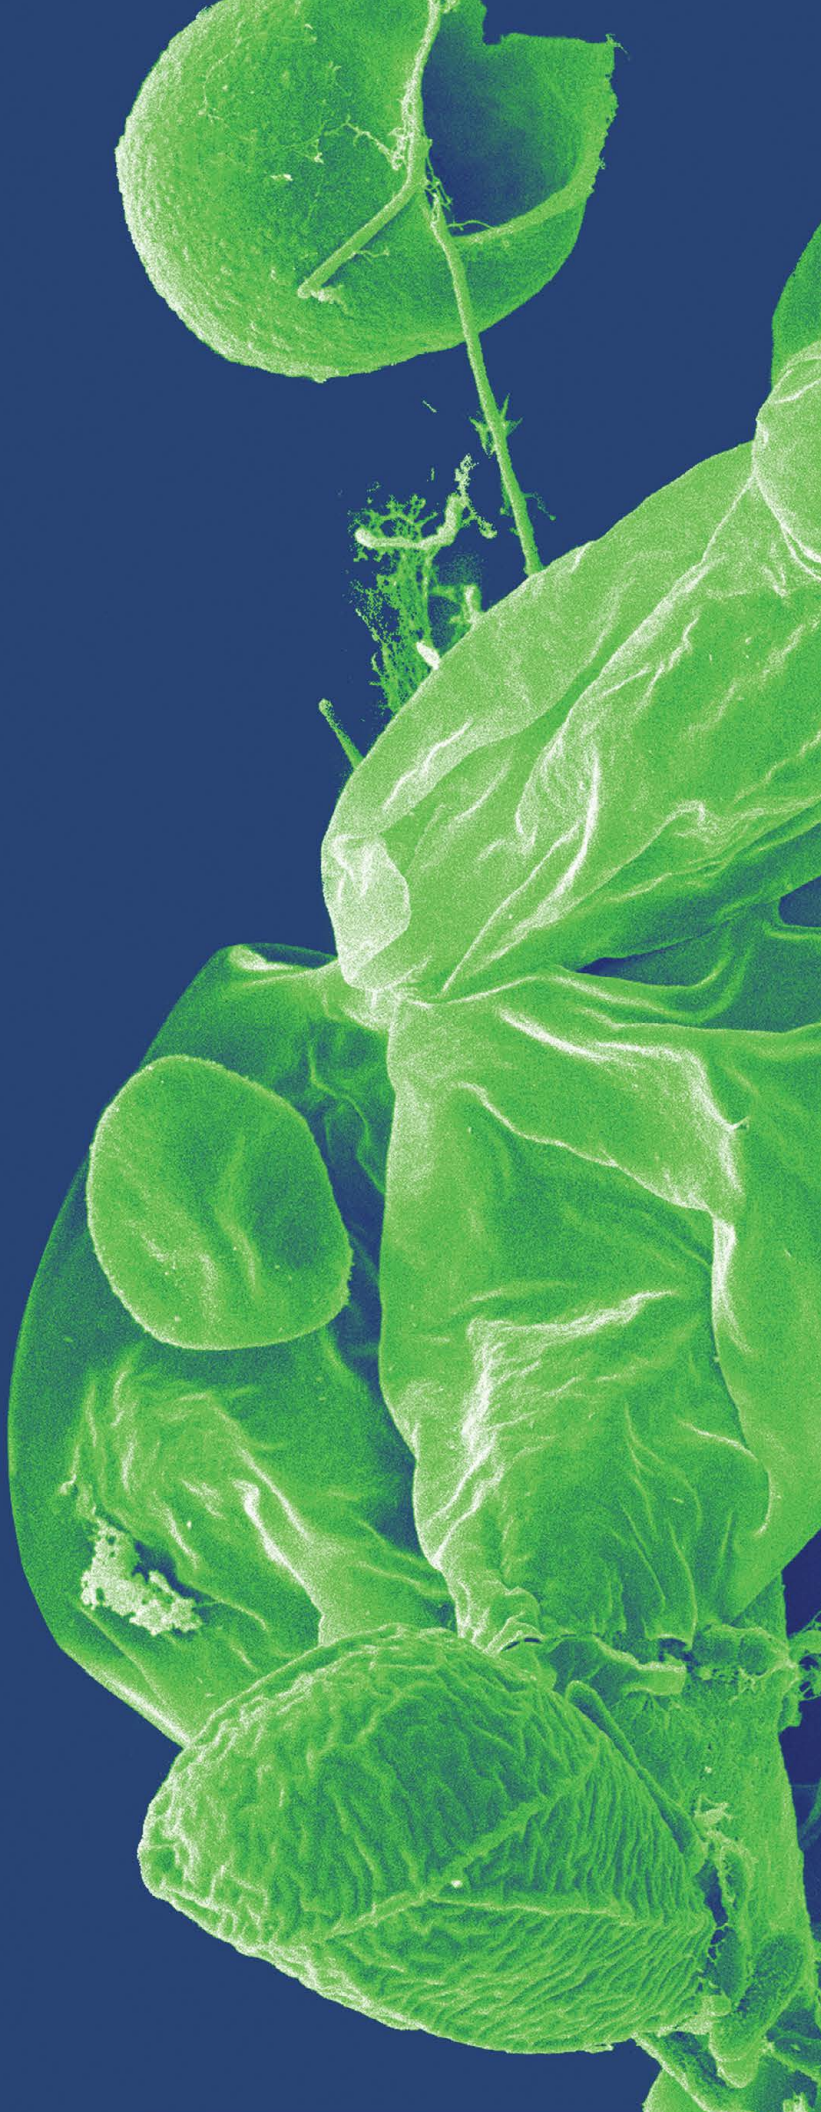

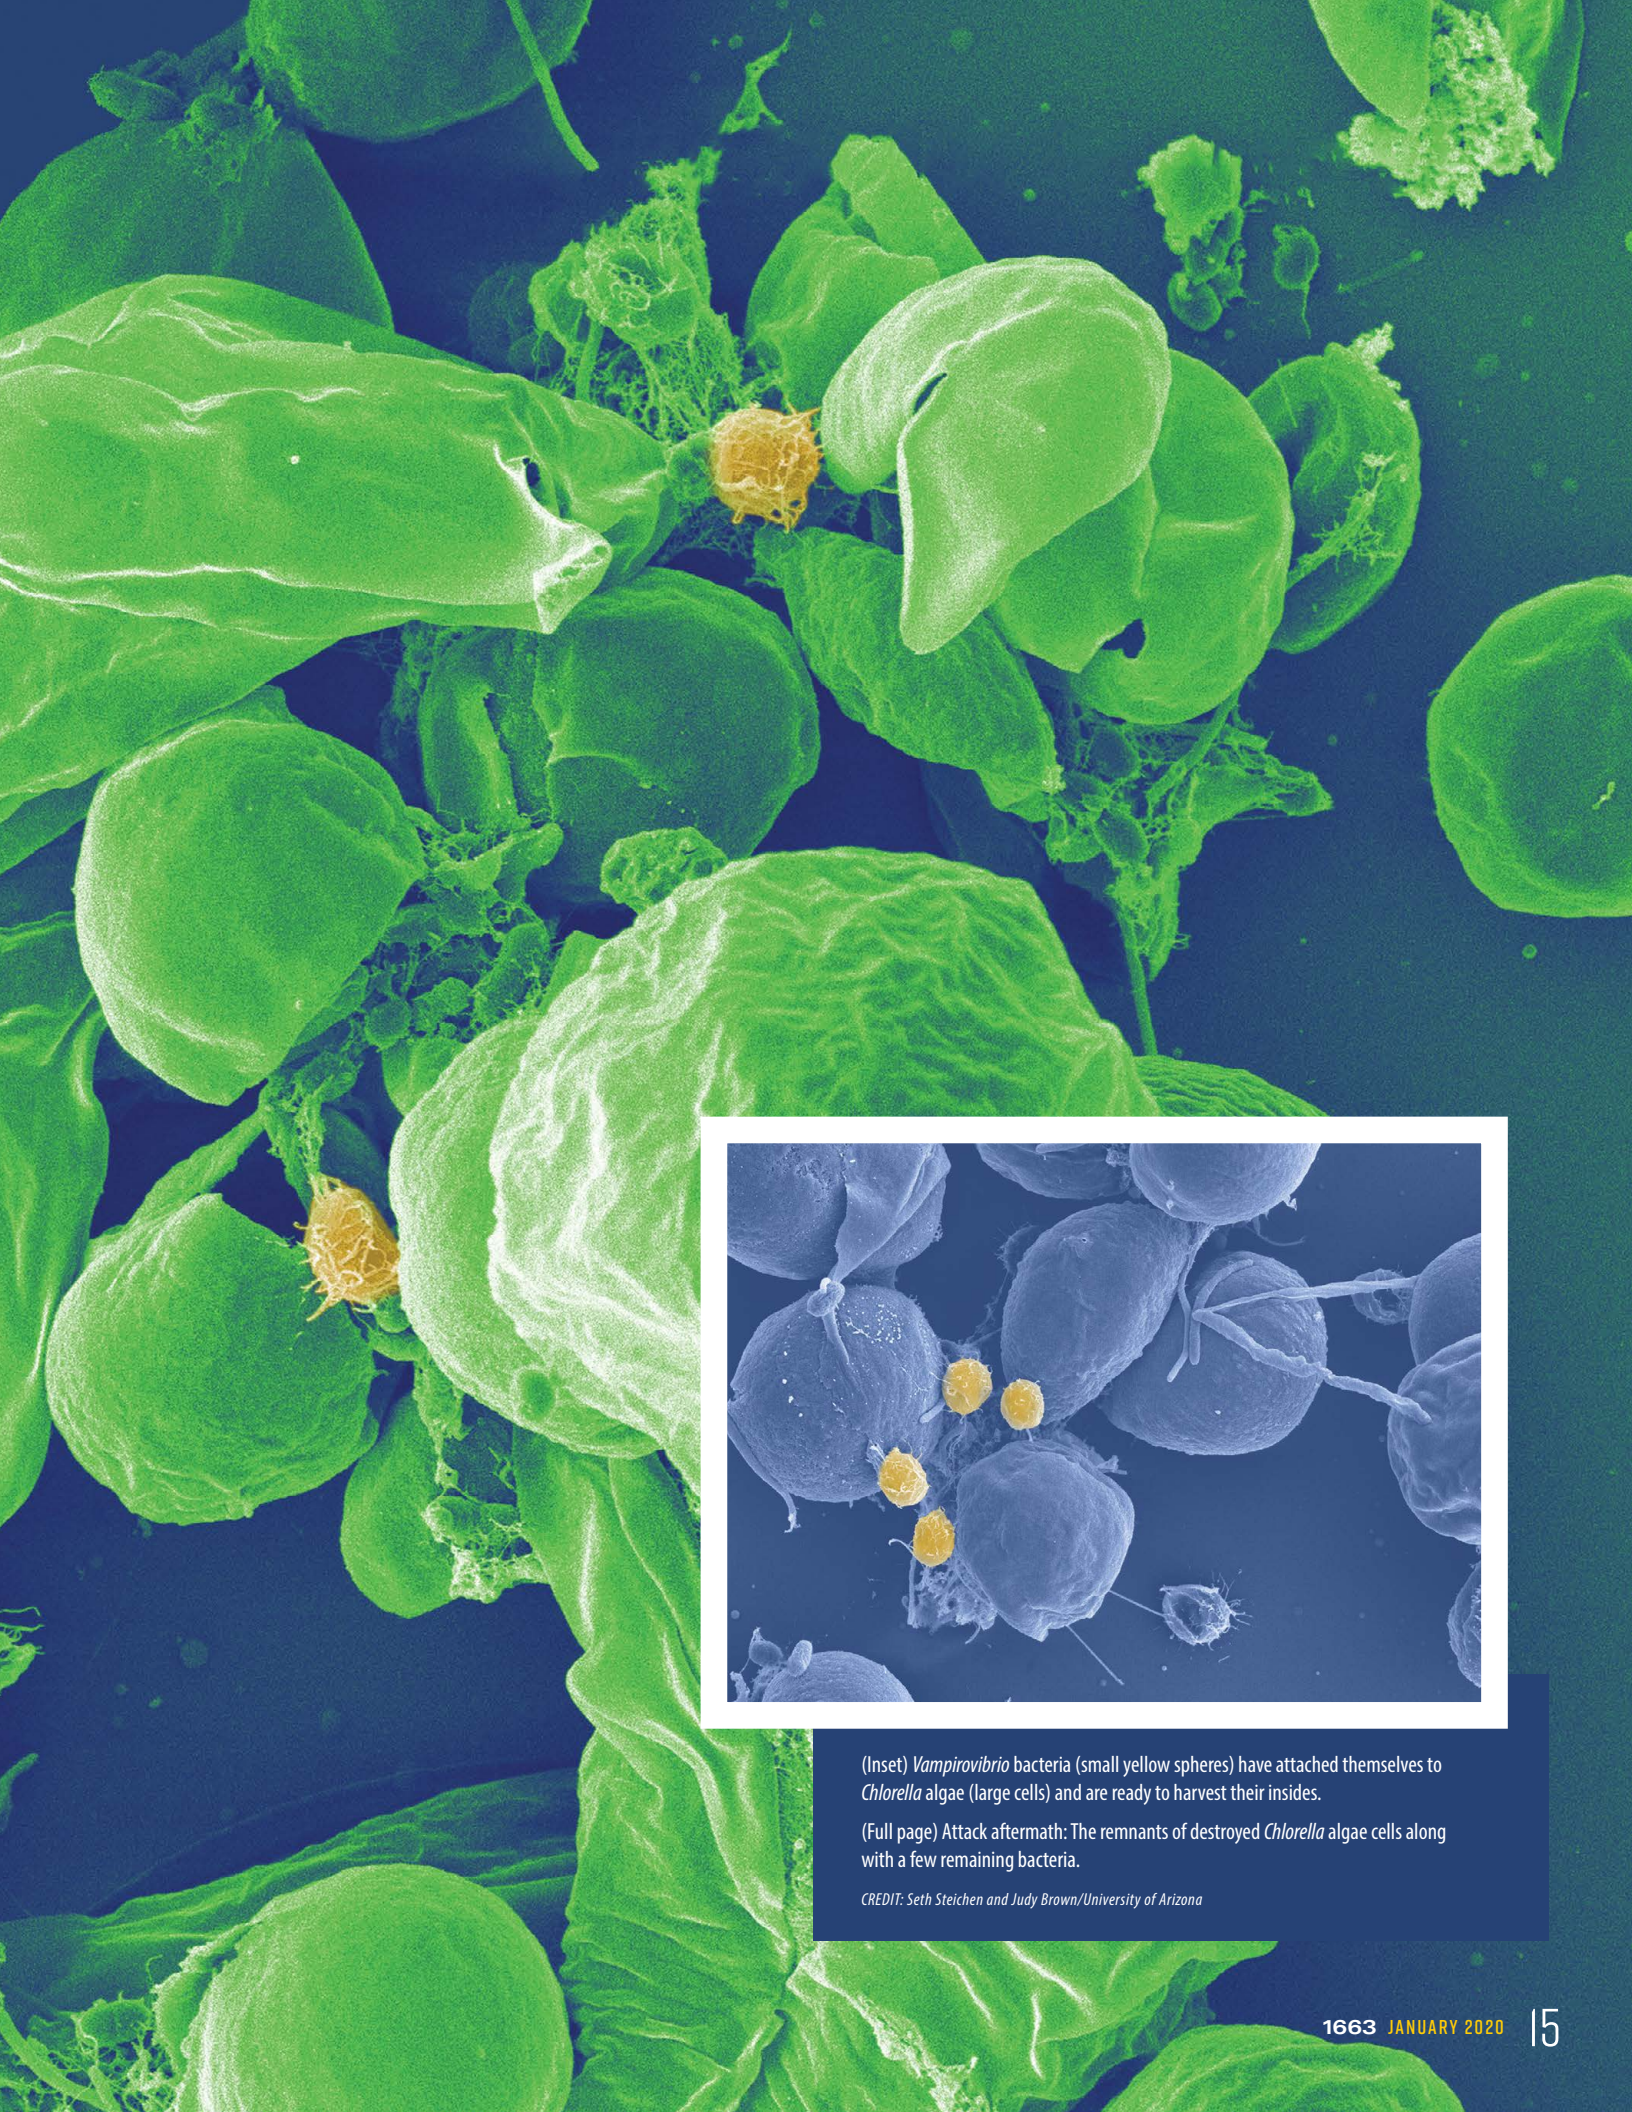

(Inset) *Vampirovibrio* bacteria (small yellow spheres) have attached themselves to *Chlorella* algae (large cells) and are ready to harvest their insides.

(Full page) Attack aftermath: The remnants of destroyed *Chlorella* algae cells along with a few remaining bacteria.

CREDIT: Seth Steichen and Judy Brown/University of Arizona

In this photograph of a Los Alamos experiment, a mirror is positioned to view molten salt inside a furnace.

REFUELING THE REACTOR

A bath of crystal-clear molten salt helps produce strategic nuclear materials and effectively transforms nuclear power into a safe, renewable energy source.

NUCLEAR ENERGY AS A RENEWABLE?

"It's not as crazy as it sounds," says Marisa Monreal, an actinide chemist at Los Alamos. "Conventional nuclear power consumes refined uranium and produces radioactive waste. Definitely not renewable. But a breeder reactor can produce more nuclear fuel than it consumes, and new designs offer extraordinary improvements in safety and energy output, with the ability to consume some radioactive waste. With fuel inputs and waste outputs so low, it comes very close to being a renewable."

Matt Jackson, a close colleague of Monreal's, is equally optimistic. "The fuel-use efficiency could be hundreds of times that of nuclear plants today," he says. "When very little fuel is wasted, very little fuel is needed."

With the right reactor design, some experts have argued, nuclear energy will deserve a place among the more widely recognized renewables, which also have limitations to their renewability. Biofuel, for example, consumes various resources, such as fertilizer, and produces waste, such as soot. Geothermal energy draws power from the heat found deep underground—heat mostly generated by natural radioactivity, a form of nuclear energy—and dredges up some toxic material in the process.

Monreal and Jackson expect that the right reactor design will see low-grade nuclear fuel immersed in a molten salt—essentially a liquefied rock that flows like water. It's just a matter of studying the bizarre stuff.

Lava laboratory

Molten salts are valuable to many nuclear technologies. Monreal and Jackson currently pursue two of these applications—power production and plutonium processing (both discussed herein)—although there are several others, such as nuclear safeguards, uranium purification, and energy storage, that can benefit from molten-salt development as well. For the sheer scope of its potential benefit, however, power production is perhaps the most tantalizing.

Monreal, Jackson, and other next-generation nuclear-power pioneers imagine eliminating the contents of today's reactors—enriched fuel rods wrapped in engineered cladding and cooled with a mechanically managed flow of pressurized water, all of which are simply entombed and replaced once spent—and instead filling a reactor with a single liquid: a molten salt, with a much more easily produced nuclear fuel simply dissolved in.

Any number of salt compounds might do the trick. One that has been experimented on extensively, for example, is a mixture of lithium and beryllium fluorides (written as FLiBe). But it can also be something more familiar, such as sodium chloride, which is ordinary table salt, or calcium chloride, another common salt. That these compounds are known as "salts" refers to their construction from positive ions (generally from the first two columns of the periodic table, like lithium or sodium) electrostatically paired with negative ions (from the second-to-last column, like fluorine or chlorine).

A mixture of chloride salts is currently loaded in an electrochemical cell in a deceptively nondescript laboratory at Los Alamos. It has been heated to more than 800 degrees Celsius, nearly 1500 Fahrenheit, causing it to melt to a liquid. But far from the thick, viscous lava one might expect from molten rock, the salt stirs easily.

"It has remarkable chemical and thermophysical properties for nuclear applications," says Jackson. "High density for desirable neutronics, low viscosity for easy pumping, lower corrosivity (compared to fluorides) for practical containment, good conductivity for catalytic processing, and high solubility for dissolving important actinide compounds into it, such as uranium or plutonium chloride."

The concept of a molten-salt reactor (MSR) itself is not new, although the push toward using an advanced fuel-salt composition is. Throughout the 1960s, Oak Ridge National Laboratory built and experimented with a small MSR. Their experiments were quite successful; the reactor worked safely and reliably, fission products and other contaminants that accumulated in the salt were generally unproblematic, and key parameters were measured that either confirmed or improved upon previous theoretical calculations. In fact, the whole enterprise seemed distinctly promising, but unfortunately, that promise came just a little too late. By the late 60s and early 70s, funding and infrastructure were already well committed to the entrenched industry of enriched solid-fuel, water-based reactors.

Today, however, interest has been renewed with proposals for potentially revolutionary improvements in both major components of the MSR: the fuel and the salt.

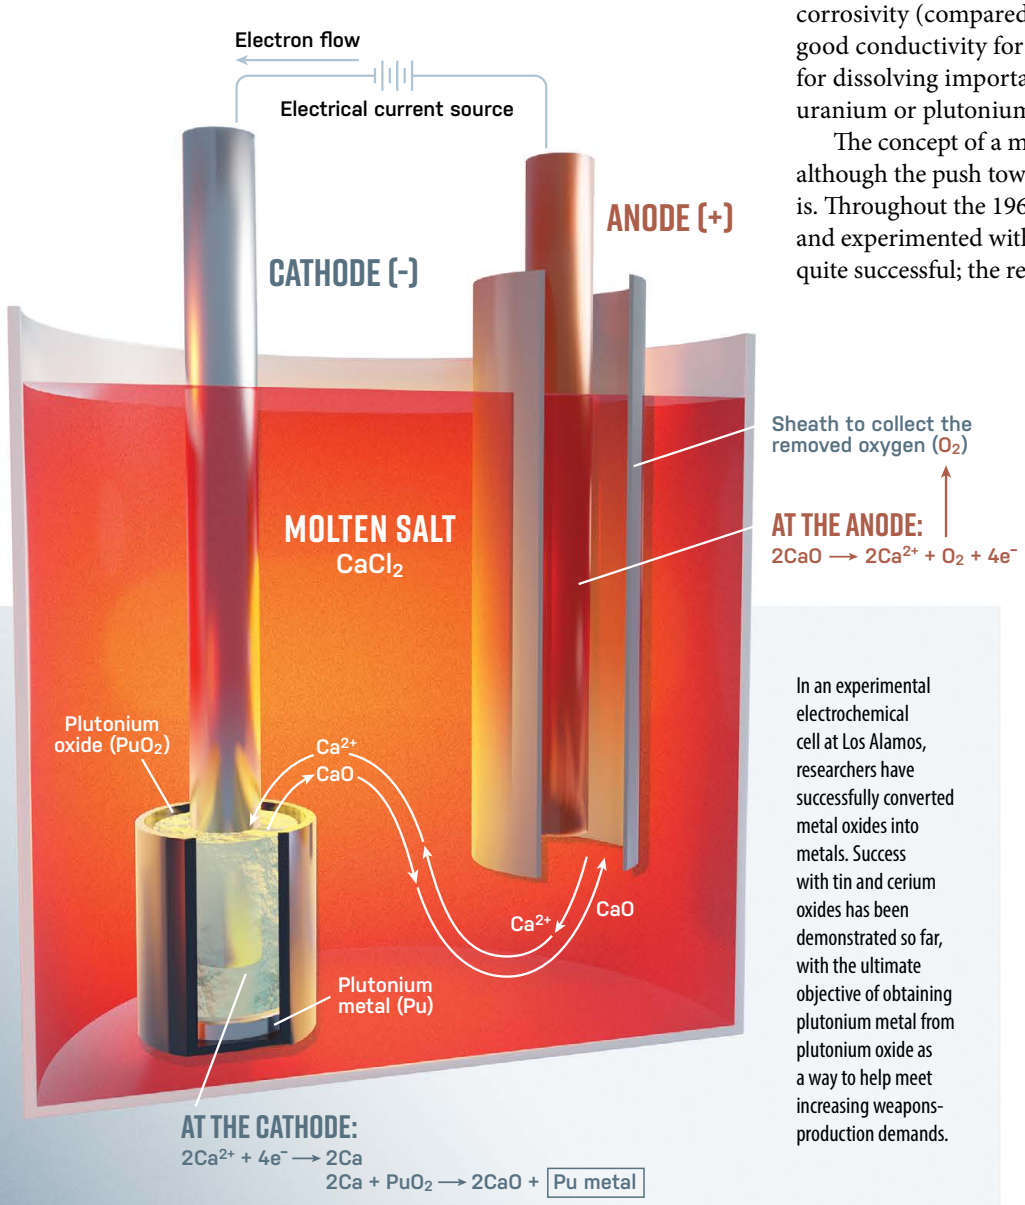

In an experimental electrochemical cell at Los Alamos, researchers have successfully converted metal oxides into metals. Success with tin and cerium oxides has been demonstrated so far, with the ultimate objective of obtaining plutonium metal from plutonium oxide as a way to help meet increasing weapons-production demands.

End of enrichment

The primary fuel tested at Oak Ridge was based on uranium-235 (U-235), which also powers most of today's commercial reactors. It comprises less than 1 percent of uranium found in nature, so it must be obtained through enrichment: separating it from the much more abundant U-238 isotope,

which itself cannot produce a sustained chain reaction to power a reactor. But enrichment is a cumbersome, costly, and wasteful process. Enriching naturally occurring uranium to just 3–4 percent U-235, which is adequate for power production, consumes a great deal of energy and typically results in less than 10 percent of the unenriched starting material becoming enriched to that level. The remainder, called "depleted uranium," gets repurposed for non-fission applications, even though it still contains a good 40 percent as much U-235 as naturally occurring uranium; this resource goes unused.

"Enrichment has been widely viewed as a necessary evil, both in today's commercial reactors and in earlier experimental ones," Jackson says.

However, just because the Oak Ridge experiment used enriched fuel doesn't mean an MSR has to. An MSR is particularly amenable to being built as a breeder reactor, which breeds its own fuel. Instead of neutrons being used to split U-235 nuclei, it can be set up for neutrons to merge with U-238. The resulting U-239 decays in short order to plutonium-239 (Pu-239), which is an excellent fission fuel, on par with U-235. The entire enrichment process can be skipped.

One need only create a uranium compound that dissolves well in the liquid salt. With a fluoride-based salt, uranium fluoride works best; likewise, for a chloride salt, it's uranium chloride. Either way, the process is not a full-blown nuclear separation but rather a comparatively simple chemical conversion, and one that is performed regularly for uranium research applications anyway. And since the uranium fuel is just dissolved into the liquid, it can remain there until converted to Pu-239 and brought to fission; it need not ever be actively removed, the way solid spent fuel is. None is wasted in enrichment, and none is wasted in the reactor. As a result, the economics of nuclear power improve dramatically. Projections show a factor of a hundred gain, possibly several hundred, in energy production over modern-day reactors per unit of unenriched fuel. As a result, very little uranium would be needed at all. One well-known estimate concludes that all human energy needs could be met for five billion years—until the sun swells up and swallows the earth—without mining any uranium ore whatsoever and instead using only a fraction of the uranium naturally found in seawater.

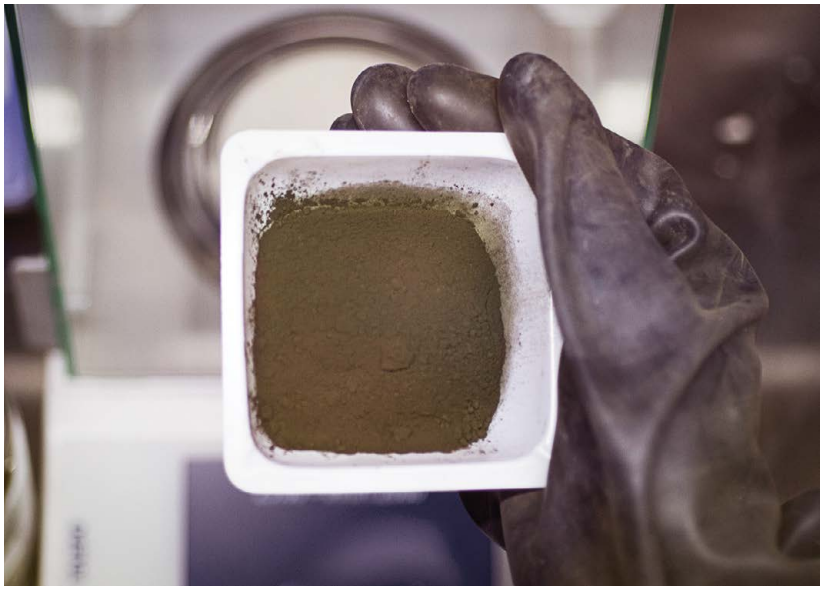

Unenriched uranium chloride: inexpensive, low-risk fuel for a molten-salt breeder reactor.

RENEWABLE NUCLEAR WOULD BE EFFECTIVELY INEXHAUSTIBLE, INHERENTLY SAFE, LOW-WASTE, NO-CARBON, ENRICHMENT-FREE POWER ON DEMAND. AND IT'S REALISTIC.

The MSR would also be vastly safer than current reactors. Meltdowns, for example, would be impossible. The fuel is already molten in normal use, and if it overheated, it would melt a plug and drain into a secure containment vessel. And it is not pressurized, so even if it were released into the environment, it would simply flow out and solidify, not explode. In addition, the public-safety risk of enriched material falling into the hands of terrorists or other enemies would be effectively eliminated. Fissile material is never manufactured, stored, or transported; rather, it is both created and consumed directly inside the reactor, where it is always dissolved in molten salt. Stealing it would be so wildly impractical as to be effectively impossible.

Pure plutonium

The Oak Ridge experiments used the FLiBe fluoride salt because it could be integrated well as both coolant and a component of their uranium-tetrafluoride fuel medium. But a chloride-based salt (and fuel) would have important benefits over a fluoride salt. Chloride salt is heavier, which improves the energy distribution of the neutrons that breed plutonium-239 and cause it to fission. It can dissolve more fuel, which means it can also accommodate a greater buildup of fission products (the resulting halves of the "split" atom) without the fuel becoming too dilute. It is lower in viscosity and therefore easier to pump through the reactor and other components. And it is safer and more economical to purify. The only real hiccup is the fact that chloride salt hasn't undergone the same extensive testing as fluoride salt. This is where Monreal and Jackson are particularly well positioned to help.

One might wonder why they, as Los Alamos scientists, would be working on a new nuclear-power reactor. While certainly known for nuclear research, Los Alamos's portfolio is weighted toward national security over commercial power. But Monreal and Jackson didn't just leap into commercial power. They started squarely in the middle of the Lab's national security mission space, developing another key molten-salt technology to address a pressing challenge.

"Los Alamos has been tasked to ramp up annual production to 30 plutonium pits—the business end of a nuclear weapon—by 2026," says Jackson. "This is a steep increase over our current capability." The ramped-up production requirements coincide with a national weapons-modernization initiative, and meeting them will not be easy; there is no simple set of knobs to turn that will generate a significant increase in weapons-grade plutonium. The effort will require new workers, new facilities, and above all, a lot of workable plutonium.

Plutonium isn't found in nature; it must be deliberately manufactured from other elements, much like the breeder MSR will do. Some can be recovered directly from decommissioned nuclear weapons, and this is indeed the primary strategy for meeting the pit-production goals. But most of the plutonium available today was produced during the Cold War and purified through a process that leaves it in oxide form: PuO_2. The trouble is, for a weapon, the oxygen must be removed and the plutonium reprocessed to obtain pure plutonium metal. Doing so by current methods is delicate, painstaking work, requiring rigorous and expensive controls to make it safe. It is also a serious production bottleneck.

Monreal and Jackson have been successfully pioneering a better way—one that will allow them to support the current pit-production challenge and expand capabilities going forward to take advantage of the larger supply of plutonium that exists in oxide form.

Due to the earth's reactive oxygen atmosphere, metals in nature are often bound to oxygen. In some cases, removing the oxygen (a process referred to as "reduction") can be accomplished with a minimal-resource, low-waste catalytic process taking place in an electrochemical cell, similar to a battery, but filled with—you guessed it—molten salt.

NEW VISION FOR FISSION

A **fission reactor** uses neutrons to split nuclei of nuclear fuel, releasing energy in the form of heat. That heat boils water to steam, and the steam turns a turbine. The rotation of the turbine drives an electromagnetic generator to produce electricity.

In nuclear power plants operating today, the reactor contains solid enriched uranium-oxide fuel rods (and separate removable control rods to quench the reaction when

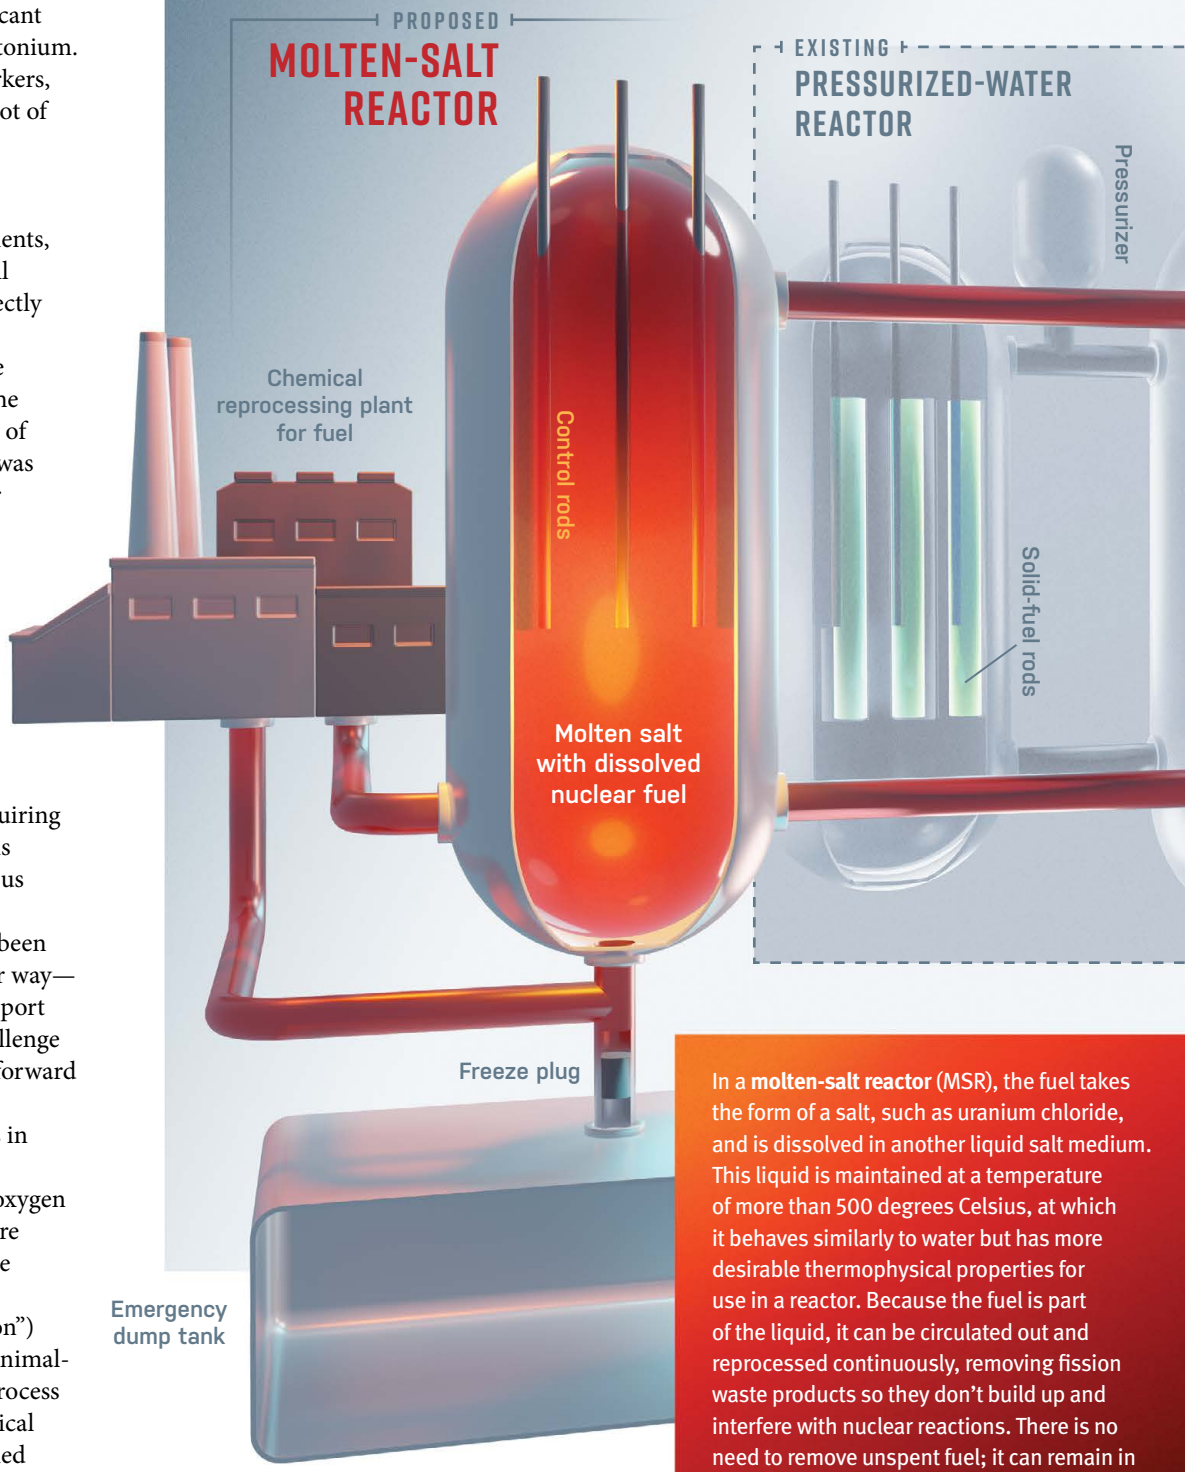

In a **molten-salt reactor** (MSR), the fuel takes the form of a salt, such as uranium chloride, and is dissolved in another liquid salt medium. This liquid is maintained at a temperature of more than 500 degrees Celsius, at which it behaves similarly to water but has more desirable thermophysical properties for use in a reactor. Because the fuel is part of the liquid, it can be circulated out and reprocessed continuously, removing fission waste products so they don't build up and interfere with nuclear reactions. There is no need to remove unspent fuel; it can remain in

needed) submerged in a fluid moderator medium, usually water, which slows down neutrons to help them induce more fissions and carries heat to the steam generator. The fuel rods become progressively less effective over time and must be replaced at regular intervals despite still containing fissile (fission power-enabling) uranium-235. The regular insertion of fresh fuel rods requires the regular enrichment of nuclear fuel—an energy-intensive, costly, and wasteful process—that increases the uranium-235 content from about 1 percent in naturally occurring uranium to about 4 percent. The other 96 percent remains non-fissile uranium-238, which ultimately becomes nuclear waste.

The reactor vessel is pressurized, and in the event of an overheating accident, such as a natural disaster that overwhelms the system's many redundant safeguards and prevents cooling water from circulating through the reactor, temperatures and pressures can rise beyond their design tolerances. Solid fuel can undergo meltdown, leading to more fission, greater overheating, and either a controlled venting or an uncontrolled explosion that disperses radioactive gases widely into the atmosphere and possibly contaminates nearby lakes or rivers that have been tapped for cooling water.

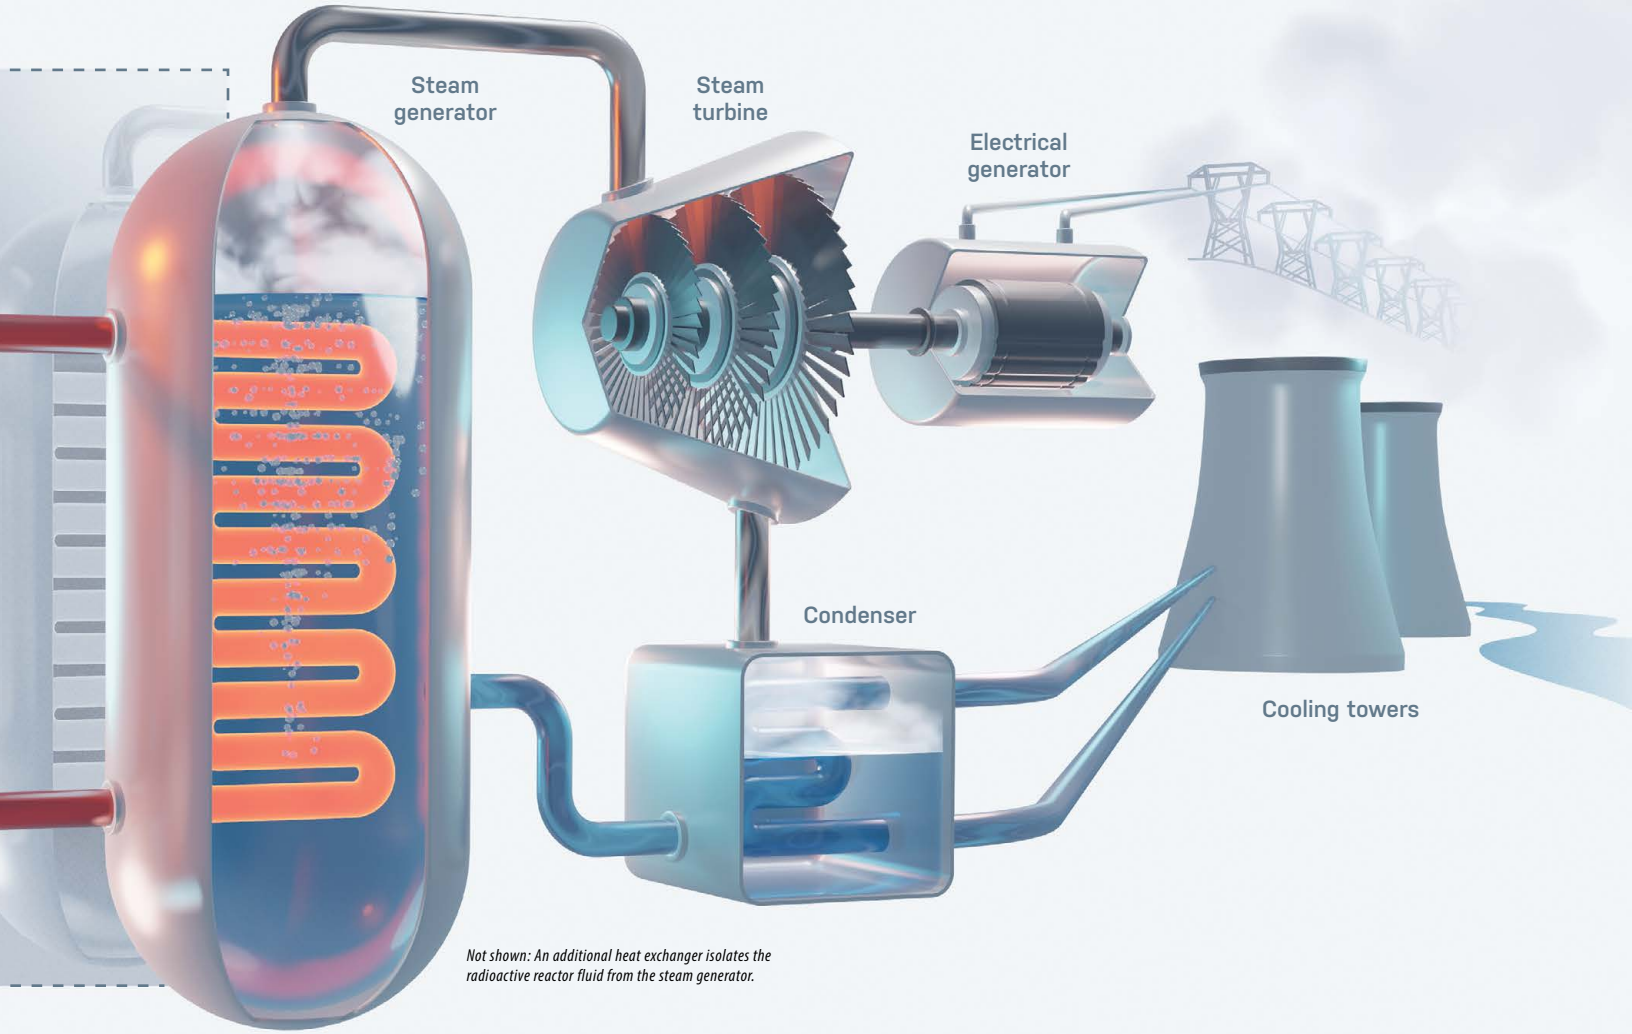

Not shown: An additional heat exchanger isolates the radioactive reactor fluid from the steam generator.

place until fission is induced, so little to none is wasted. The energy output per fuel input may be hundreds of times greater than that in existing nuclear reactors.

Molten salt is a poor moderator. An additional moderator can be added, or the neutrons can remain unmoderated, zipping about at high speed. This reduces the fission rate but increases the rate at which non-fissile uranium-238, which makes up more than 99 percent of naturally occurring uranium on Earth, is converted to fissile plutonium-239. Such a molten-salt fast reactor (MSFR) is also known as a breeder reactor, in that it produces, or breeds, fissile fuel faster than it consumes that fuel. In this sense, an MSFR borders on being a renewable-energy system. In addition, it runs on a small amount of uranium-238, eliminating the need for uranium enrichment entirely, and it can consume certain dangerous isotopes during normal operation, thereby downgrading nuclear waste. Furthermore, because the fissile fuel is produced and consumed entirely inside the reactor, it cannot be stolen, whereas enriched solid fuel can be stolen at the enrichment facility, the power plant, or in transit between the two.

A nuclear accident in any type of MSR, or even deliberate sabotage, would not produce a meltdown because the fuel is molten already. In the event of overheating, the molten salt, which is not pressurized, melts a plug and drains harmlessly into an underground containment vessel. This occurs naturally without any operator actions, mechanical systems, or computer controls, so there is no danger of radioactive release. And even if an earthquake or similar low-probability disaster were to produce a release, the molten salt would solidify in place rather than disperse into the environment.

The process had already been discovered in academia and successfully employed and scaled up for the production of titanium. Monreal and Jackson suspected they could make it work for plutonium.

Essential electrolytes

Of course, even at Los Alamos, one of the few places capable of working with plutonium, one does not simply grab plutonium oxide off the shelf and begin experimenting with it. So Monreal and Jackson, joined by Los Alamos engineering colleague Kirk Weisbrod, set up a research plan to test surrogate metal oxides first. Weisbrod was instrumental in designing and building a suitable electrochemical cell—a molten salt-filled vessel with two electrodes and a current-carrying external circuit in between. First, they tested it, successfully, with tin oxide. Then they tried cerium oxide because cerium is one step closer to matching the electrochemical properties of plutonium. This was also successful, but there were complications because of several stable, partially reduced cerium compounds, such as CeOCl and Ce_2O_3, to pass through along the way from CeO_2 to cerium metal. The final step, currently underway in a collaboration with other interested laboratories, is testing the process with plutonium oxide. This should actually be easier than cerium oxide, since the reduction mechanism for plutonium oxide mitigates problematic partially reduced compounds. Additionally, plutonium oxide is conductive at higher temperatures, therefore improving overall efficiency.

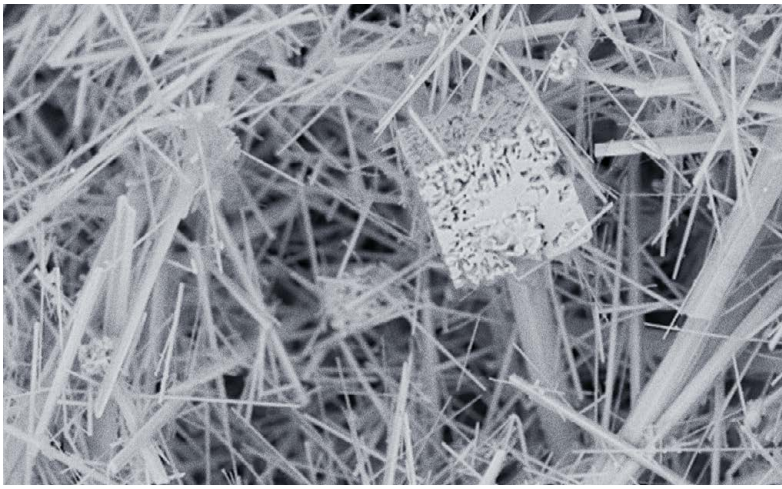

Scanning electron microscope image of $CaSnO_3$ (a calcium-tin-oxide, or calcium stannate) formed on the anode during a molten-salt electrochemical experiment intended to strip the oxygen from tin oxide in preparation for a similar experiment with plutonium oxide. These crystals—both the needles and the cubes come from this same compound—are electrically insulating; as they built up, the electrochemical cell lost functionality. The problem can be solved with a different anode material or a different metal oxide.
CREDIT: Scott Parker/LANL

ISOLATING PLUTONIUM METAL BY CURRENT METHODS IS COMPLICATED WORK— AND A SERIOUS PRODUCTION BOTTLENECK.

Molten salt already plays a role in Los Alamos's existing process for reducing plutonium oxide to plutonium metal, but that is a purely chemical process, not an electrolytic one. In this existing chemical process, the PuO_2 and granulated calcium are mixed into molten calcium chloride salt, which is used here as both a solvent and a heat sink. The calcium displaces the oxygen and liberates plutonium, which settles to the bottom. Then chlorine gas is added in order to convert the excess calcium and the calcium-oxide byproduct (containing the oxygen removed from the plutonium) back to calcium chloride; upon chlorination, the oxygen is allowed to escape.

It may sound straightforward, but there is a three-fold difficulty. First, both added ingredients, calcium and chlorine, are hazardous materials, requiring safety precautions and monitoring by trained personnel. Pure calcium is explosively reactive in water, and chlorine gas is poisonous. Second, even though the process has been greatly optimized, it still requires these two separate steps—adding calcium to remove the oxygen, adding chlorine to restore the salt—which must be repeated four times to obtain the plutonium metal in a typical batch. This is quite time-consuming. Third, and also time-consuming, is a separate purification step needed at the end. Contaminants are removed from the plutonium metal through a process called electrorefining, which involves setting up a current in a separate electrochemical cell. When all is said and done, the molten salt is no longer usable and is disposed of as radioactive waste.

In the new process, the same calcium chloride salt is used, this time as an electrochemical conductor. The negatively charged electrode, or cathode, catalyzes the key swap: separating oxygen from the plutonium and attaching it instead to calcium ions in the salt. The calcium oxide thus produced then finds its way to the positively charged anode, where the oxygen is stripped from the calcium oxide and removed, restoring calcium ions to the salt. That reaction produces free electrons, which are then carried as a current along an external circuit to the cathode, restoring the negative charge there to continue the process. Neither granulated calcium nor chlorine gas is needed, removing the hazards so that the process may run unattended, and the salt is restored for future reuse, rather than discarded. Furthermore,

Monreal and Jackson are hopeful that since both the reduction and purification processes are electrochemical in nature, they can be combined in the same electrochemical cell. All in all, the total processing time can be significantly decreased.

Salt substitute

Now that Jackson and Monreal have built and tested their electrochemical cell for plutonium reduction and purification, it is available as a testbed for other experiments—in particular, for research that might facilitate the adoption of essentially renewable nuclear power. Ultimate success would offer effectively inexhaustible, inherently safe, low-waste, no-carbon, enrichment-free, ample power on demand, whether it's sunny or windy or neither. And it's realistic.

What does MSR-based nuclear power need to become a reality?

"A lot of things, but none of them is prohibitive," says Monreal. "We just need to further develop our facility for working with these molten chlorides. We need to continue to measure them, model them, control them. We need to develop the industry to go along with the reactor."

The industry Monreal mentions will be one of sensors, systems, software, and safeguards. None of this is easy; even just the sensors—which will be needed to monitor the fluid properties, contaminant levels, and corrosivity of the salt solution in the reactor—do not exist. How do you reliably and continuously measure the properties of molten lava? Even for something as simple as its volume, which is a key component of density, it is not immediately obvious how to measure it inside an opaque, sealed furnace, using only components that can withstand the heat.

"I think we can design a new kind of dilatometer to handle that," says Jackson, referring to an instrument that will sense the expansion of the salt as it is heated. "But that's just volume. Other measurements will be harder."

One of these harder measurements is fluid viscosity—say, at different temperatures or with different levels of actinide fuels or fission products mixed in. It is an important parameter to understand for MSRs, but the most direct method—dropping spheres of different densities through the medium and recording their fall time—requires seeing inside the furnace. Jackson and Monreal are using an advanced Los Alamos capability, neutron imaging, to see *through* the furnace instead. The research plan is to conduct this experiment repeatedly in a laboratory setting—experimentation that's already underway and producing results—in order to tabulate the range of viscosity and density values for the salt at different temperatures and actinide concentrations. Such data would then feed into new simulation software for predicting the molten material's behavior as conditions are changed, for plant operation or optimization purposes.

Because the medium will support nuclear reactions, it will be creating different elements as it runs—breeding plutonium, splitting into fission products, absorbing neutrons, radioactively decaying from one element or isotope to another—and any or all of these changes could result in markedly different properties, which is why Jackson and Monreal are experimenting with different compositions and concentrations. Some might alter the overall heat capacity or conductivity. Some might spawn the production of corrosive agents. Some might require specific types of chemical analysis and processing alongside the reactor. Some might imply a tradeoff: do X and you'll get better Y but worse Z. To make MSR-based renewable nuclear power a reality, all of this must be known.

WHEN VERY LITTLE FUEL IS WASTED, VERY LITTLE FUEL IS NEEDED.

And all of it will be.

The United States currently gets about 20 percent of its electricity from 97 nuclear reactors with an average age of nearly 40 years. All of them run on enriched solid fuel, none is a breeder, and none is anywhere near as safe, efficient, or environmentally friendly as an MSR. If it sounds like the time is right for a nuclear renaissance, the experts agree.

"We could see molten-salt reactors ready for commercial construction in the near future," says Monreal, "once we master molten salt." LDRD

—*Craig Tyler*

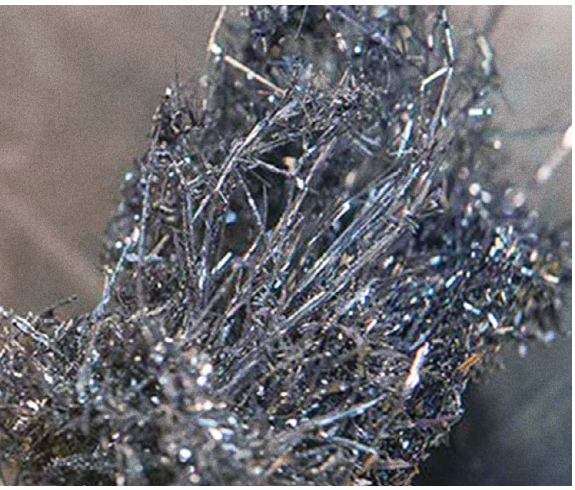

Using a modified configuration of the electrochemical cell, these high-purity uranium dendrites were produced from impure uranium metal in a molten-salt electrolyte. This purification step comes after the conversion from metal oxide to metal.

CREDIT: Matt Jackson/LANL

MORE **NUCLEAR INNOVATION** AT LOS ALAMOS

http://www.lanl.gov/discover/publications/1663/archive.php

- **Working with nuclear materials**
  "Resource Revolutions" | July 2019
  "Intelligent Infrastructure" | February 2019
- **Small fission power plants**
  "Megapower" | February 2019
  "Power to the Planet" | August 2018
- **Fusion power research**
  "Mission: Ignition" | February 2019
  "Small Fusion Could Be Huge" | July 2016
- **Subcritical weapons experiments**
  "The Bomb without the Boom" | October 2017

from 1663

The Mind in the Machine

Machine learning is a **scientific revolution** that is changing how science gets done.

WE'RE NOT TALKING ABOUT ROBOTS (though it would be cool if we were). Sentient humanoid automatons, whether benevolent or malevolent, walking and talking amongst us, are still science fiction. But some things that used to be impossible are now science fact—like computers that can tell if two disparate images are actually showing the same thing or that can predict if and when a supercomputer will crash. Los Alamos has always excelled at data science, and the data-science techniques known collectively as machine learning are now taking data analysis to the next level. Through machine learning, or ML, scientists are exploring new ways of answering old questions, and, in some instances for the first time, they are actually getting some answers.

Machine learning is a natural product of increased computational power. The questions aren't necessarily new, and the math isn't necessarily new. But the machines are, and what scientists are doing with them certainly is. Enabled by major advances in computer hardware and software, and by the massive amounts of data newly available, tech entities from social media companies to national laboratories are using and developing ML.

But while social media and computer companies are mostly working on problems like targeted marketing, virtual assistants, and self-driving cars, Los Alamos scientists are working on mission-critical science problems like nuclear nonproliferation, global security, and ensuring the safety, efficacy, and reliability of the country's nuclear arsenal. The level of performance required for the Laboratory is more stringent owing to the high-stakes nature of these challenges. What sets the Laboratory apart from other entities pushing ML is the intersection of problems and solutions found here—Los Alamos offers leading-edge scientific solutions rooted in rich institutional knowledge. The broad body of physics expertise that exists at the Laboratory, when married to ML, makes new approaches to national security possible.

Scientists at the Laboratory are using and developing ML in a plethora of ways. Some are going after answers to age-old questions, some are asking brand new questions, and some are pioneering new ways of doing and thinking about ML itself. By no means comprehensive, this article provides several examples of machine learning being done at Los Alamos.

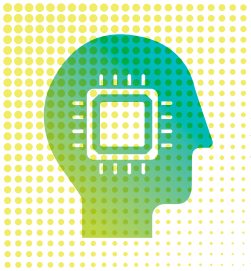

What it is and what it isn't

ML is not synonymous with artificial intelligence (AI). General AI refers to learning and reasoning by machines without the intervention of humans, and most scientists agree that we aren't there yet. ML is a specialized subset of AI, wherein a human still writes the code, but the output of the code depends on data—usually vast amounts of it—that is also chosen and fed to the computer by a human. And the computer usually needs to be told, by a human, whether or not it is doing the right thing with the data.

Machine learning is a natural product of increased computational power.

The U.S. Defense Advanced Research Projects Agency describes the evolution of AI thus far as having occurred in three distinct waves: The first wave, from the 1970s through the 1990s, was characterized by computers with the ability to reason—but not learn or generalize—as illustrated by IBM's Deep Blue, the chess-playing computer that repeatedly beat the reigning human world champion. The second wave, from the 2000s through the present, is characterized by computers with the ability to learn and perceive—but not to generalize—as illustrated by virtual assistants like Apple's Siri and Microsoft's Cortana (and others). The third wave, mostly still in research labs and not yet producing products ready for mass consumption, will likely last for 10–20 years and be characterized by computers that can reason, learn, perceive,

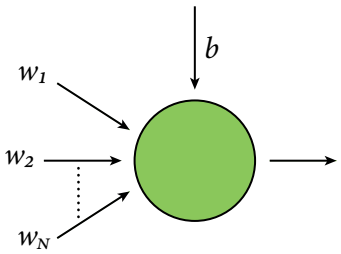

In a neural network the nodes are called neurons, and each neuron has learnable parameters: multipliers (*w*) are called weights, and addends (*b*) are called biases. The exact values of the weights and biases for each neuron are arbitrary at first and become progressively more finely tuned through the iterative training process.

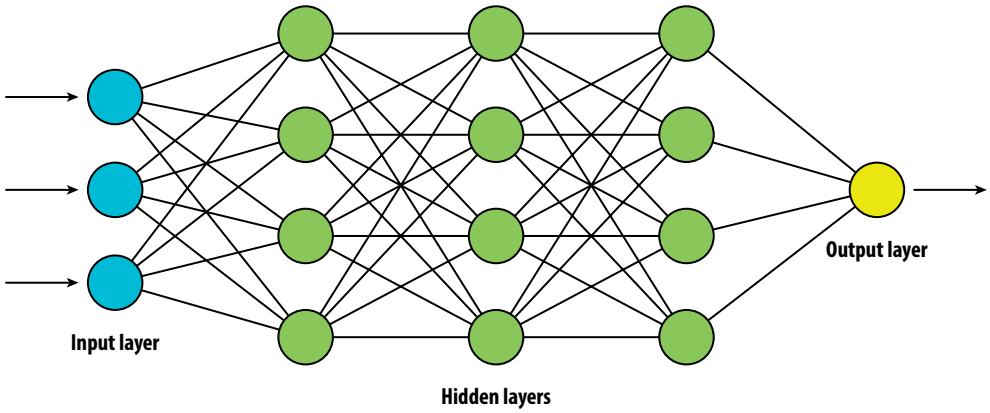

A neural network is a common machine learning method designed and named after biological neuronal systems. The mathematical computation occurs in the hidden layers, which may consist of any number of neurons from one to hundreds. Similarly, the network itself may have anywhere from one to hundreds of hidden layers between the input and output layers.

and generalize. A third-wave computer would, for example, be able to converse in natural language and explain to a human its decision-making process.

So ML is not quite AI, but nor is it simply good programming. For example, mobile navigation apps that find the fastest travel routes aren't using ML; they've just been intelligently programmed. For it to be ML, the machine has to learn an algorithm without being explicitly told how. It's a bit like human toddlers learning to walk. They practice a little and learn what works, or even hold a parent's hand, but the parent isn't saying, "transfer your weight to your right leg; bend your left hip, knee, and ankle simultaneously to lift your left foot off the ground while keeping your right leg straight; move your left leg forward several inches before placing your foot, heel first, back on the ground." And even if the parent *were* saying all that, those instructions would not compute in the toddler's brain. The parent says, "this is walking, now you walk," and the child figures it out. The steps, or algorithm, are internally generated—they aren't specified by the external programmer (parent), and when the desired outcome is achieved, the machine (toddler) has learned.

With traditional computer programming, data and rules go in, and answers come out. With ML, however, there are two phases: during the training phase, data and answers go in and rules come out; during the inference phase, data goes in and predictions come out.

One popular ML method, out of many, is the use of a neural network. Named after the way neurons organize in brains, an ML neural network is a mathematical model that is organized according to an extreme simplification of living neural systems. Neural networks can be simple, consisting of an input layer, output layer, and single computing layer in between, or they can be complex, or "deep," with many computing layers in between the input and output layers. The computing layers are called "hidden layers" because the user doesn't interact with them, as with the input and output layers. As a calculation progresses through the layers, the resolution becomes finer. Each layer of a DNN can learn a different concept, so part of the challenge for scientists is figuring out the best combination and order of layers.

Training a deep neural network (DNN) is just minimizing a cost function using a collection of known inputs and outputs called a training dataset. Each input goes in, gets worked on by the succession of hidden layers, then emerges as some output, which is at first not very close to the expected known output. The difference between expected output and actual output is the cost, and the way to reduce it is to go back to each neuron, or node within a hidden layer, and adjust what it does, mathematically, to the input it receives. Every neuron has "weights," by which inputs get multiplied, and a "bias" that gets added to the output. The weights and biases are the learnable parameters, so their values are arbitrary at first and get repeatedly modified through the iterative training process. Once the cost function is minimized, meaning all the weights and biases are dialed in so that the actual output is as close to the expected output as it can get, the model is trained.

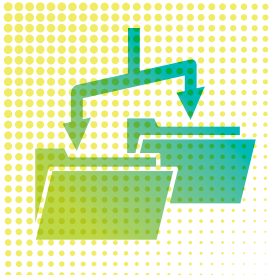

Graphics and schematics

One thing that brains have historically done better than computers is vision. But brains are helping the computers get better. Online reverse-image search tools, for example, when provided with a photo, can find other instances of the same photo or other visually similar photos. And when it comes to classification—is this a photo of a dog or a cat?—ML models are now performing better than humans. But what if the image of interest is not a photograph, but a technical drawing, like a wiring schematic, blueprint, or data graph? In those instances, brains still reign supreme.

For an ML model to assign a photograph to a predefined class (e.g., dog vs. cat), the model compares numerical values for each pixel to the values of its eight neighboring pixels. But unlike photographs, which contain information about color and intensity in nearly every pixel, resulting in texture and shading throughout the image, technical drawings are line drawings that have very little per-pixel information, so standard image-classification ML won't work. Los Alamos ML expert Diane Oyen is working on ML models to automate the analysis of technical drawings for the Laboratory's nuclear counter-proliferation mission.

"Just like ML models for classifying photographs won't work for technical drawings, our ML approach for technical drawings wouldn't work for photos," says Oyen. "Diagrams have a lot of white space, so our models don't have to go pixel by pixel."

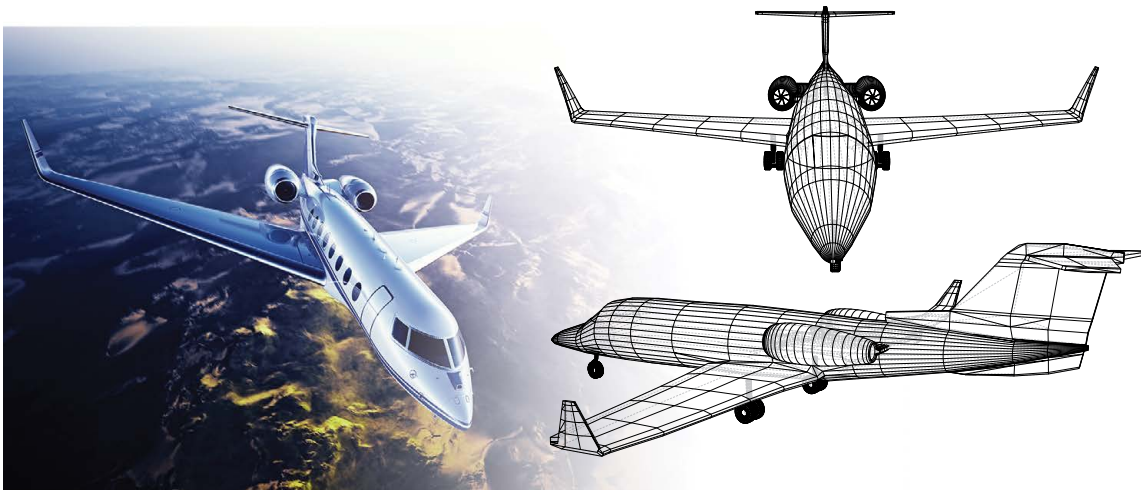

Most image-classification ML models work well on images like this one of a jet, but they fail miserably on technical drawings like these schematics for a similar aircraft. To a typical adult human brain, it is almost immediately apparent that the two technical drawings show the same item from two different angles. Diane Oyen is developing a computer model that will similarly be able to determine that the two drawings show the same thing, despite the fact that they have very few lines in common.

ML photo classifiers fail for technical drawings because these models don't retain spatial relationships—a circle inside a square is the same as a circle next to a square. But Oyen and colleague Liping Yang use a graph-based approach that preserves these relationships. A graph, in this context, is a logical structure defined as an ordered pair of a set of vertices and a set of edges. How images get represented graphically is to first define points, edges, and the relationships between them, then to organize and store this information in a data matrix. Oyen and Yang's method uses non-ML classical computer-vision techniques to delineate the edges and corners of a technical drawing, then a graph-based approach to extract the topology—meaningful spatial relationships within the image. In the graph-based approach, different model layers do different tasks for organizing topological information: some layers combine lines and surfaces into shapes, such as circles and squares, and other layers describe the spatial relationships among those shapes. Then the extracted topological features are used to train a standard DNN to image-match the drawing against a collection of known drawings.

Because there is so much ML in development, a developer doesn't always have to train a new model—he or she can use pre-trained models. Taking pre-trained models and stringing them together in new ways, or modifying them for a new purpose is called transfer learning and is one of Oyen's specialties. The main benefit of transfer learning is that it shaves off a lot of the data and computation needed to train the model. DNNs can have millions of weights and biases, but those contained in the early layers of the model (just after input) are very general and affect the rough approximation. Therefore only the later layers (just before output), where more complex features get resolved, need to be changed to adapt the model for a new use.

Transfer learning is particularly useful for the kind of datasets found at the Laboratory. Many ML models get trained with internet-scale data, consisting of millions of labeled examples (e.g., photos of dogs and cats). But Los Alamos datasets often consist of a small number of highly specialized examples, and labeling them requires a human expert and a lot of time. Being able to transfer what a model has learned on a large, less-specialized dataset to a small, highly specialized dataset is invaluable.

So far, Oyen's model is good at matching replicated images, which, Oyen says, has uses beyond national security applications, such as detecting plagiarism. In scientific research, careers are built on the novelty and ingenuity of data and designs, so plagiarism is a high-stakes affair. If, say, a scientist publishes a data chart in a research paper, then a screenshot of that chart is used in a presentation slide by another scientist, and someone photographs the slide during the presentation and posts that photo to a social media page, a tool like Oyen's could connect the social media photo to the original research publication and the image's rightful owner.

Tech entities from social media companies to national laboratories are using and developing machine learning.

The way to secure ownership over a new idea or design is to obtain a patent. Here too, Oyen's image-classification tool will be of use. If two different scientists have invented highly similar items, it's unlikely that their individual technical drawings will be identical. Differences in perspective alone would be enough to confound most image-matching computer programs, so the images all have to be evaluated by human brains. Pairwise comparison of hundreds of thousands of images quickly becomes mind-numbingly tedious. Oyen is working to bring semantic information into her technical-drawing-classification tool, so that it will be able to confirm that two images, which contain the same shapes in the same relationship to one another despite differences in scale, rotation, occlusion, or perspective, are indeed showing the same thing.

Through transfer learning, the model can be adapted for other purposes. It can be customized so that subject-matter experts can use it without needing an ML expert on hand. This is called interactive learning, which is a subfield within transfer learning. The model is created from a preexisting model (transfer), but it's built in such a way that the end user can modify it (interactive) according to his or her own needs.

"There are two ways to capture domain knowledge, or subject-matter expertise, in an ML model," explains Reid Porter, Los Alamos interactive learning expert and a colleague of Oyen. "You can either go in and talk to a domain expert, then build an ML tool specific to that person's needs, or you can put the ML directly into the domain expert's hands and let him or her customize it by using it."

Interactive learning is useful for highly specialized applications because it enables general-purpose tools to be fine-tuned on the data at hand, which is often limited in quantity. Microscope

images, satellite images, and time-series data are examples of limited, specialized datasets that would be candidates for analysis by an interactive learning approach. For a real-world, nuclear nonproliferation example, consider the scenario of federal agents encountering some sort of radioactive material being transported across a border. Samples must be sent to domain experts at Los Alamos for determination of the material's provenance. Every sample is unique, so the data are limited, yet confidence is crucial. With an interactive ML tool, experts could automate some parts of, say, microstructural image analysis, in order to spend more of their time on non-automatable conclusions and validation.

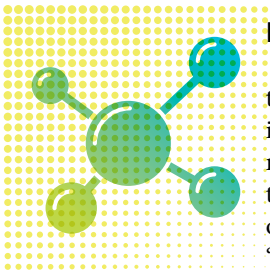

Faster output

From its inception over 70 years ago, the Laboratory has been a world leader in computer simulations of atomistic and molecular systems. ("Atomistic" refers to the tracking of each atom in a collection of atoms, as in a material, in contrast with "atomic," which generally refers to single atoms and their substructure.) Predicting how groups of atoms or molecules will interact with one another is and has always been central to Los Alamos's mission.

Laboratory physicist Nicholas Lubbers develops physics-informed ML methods to help materials scientists, chemists, and molecular biologists model the chemical and physical properties of the atoms and molecules they study. They want to know, for example, how a shock wave will propagate through a certain kind of metal, or how an enzyme will interact with DNA inside a human cell. For these domains, a model has to obey the laws of physics, but mainstream DNNs don't typically include even basic physics principles. Lubbers and his colleagues have been working to marry the flexibility of ML techniques with the constraints of physics. They build DNN architectures that encode exact physics properties, such as translation and rotation invariance, as well as approximate properties that are found in atomistic systems, such as locality. The result is physically valid ML models that are much more robust and accurate for large, complex systems.

When it comes to methods for computationally modeling the energies of individual atoms, bonds between atoms, clusters of atoms, individual molecules, and molecular interactions, there tends to be a tradeoff between affordability and accuracy. Methods using the equations of classical physics are very affordable but are of limited accuracy and lack transferability to other domains, while methods using the equations of quantum physics are highly accurate and transferable but scale poorly—quickly becoming cost restrictive as the number of atoms in a simulation grows.

The scientists have been using transfer-learning techniques to develop best-of-both-worlds solutions. Basically, they train a DNN on vast quantities of classically computed approximate data to get the model's basic structure, then retrain the last few layers on higher-quality but lower-quantity data generated by the best quantum-physics calculations to perfect and polish the model. In this way, the model operates with classical cost and quantum accuracy, allowing scientists to simulate larger systems over longer time scales. If necessary, the ML model can occasionally be compared to quantum-mechanical calculations as a sort of spot check to ensure accuracy.

When training an ML model, it is possible to overtrain, which undermines transferability. For example, if you were

ML is not quite AI but it is much more than good programming.

Nicholas Lubbers and colleagues use ML to accelerate the simulation of various complex systems. This simulation shows an experimental drug molecule interacting with a protein from *Mycobacterium tuberculosis*, the pathogen that causes tuberculosis. Some bonds between atoms are flexible, allowing the drug molecule to take different shapes, each of which might interact differently with the pathogen protein. Lubbers' model predicts how much energy it takes for each possible shape to form, which tells the drug designers how likely it is that the drug will take that shape when it interacts with the pathogen protein. For example, as the right-hand portion of the drug molecule rotates about the indicated carbon-nitrogen bond (torsion angles from −180 through +180), thereby changing the molecule's shape, the energy of that bond also changes. Every atom of the pathogen protein and the drug molecule is included in the simulation, which helps drug developers understand how the molecules might be expected to interact, thereby helping them assess how effective the drug might be. *CREDIT: Justin Smith/LANL*

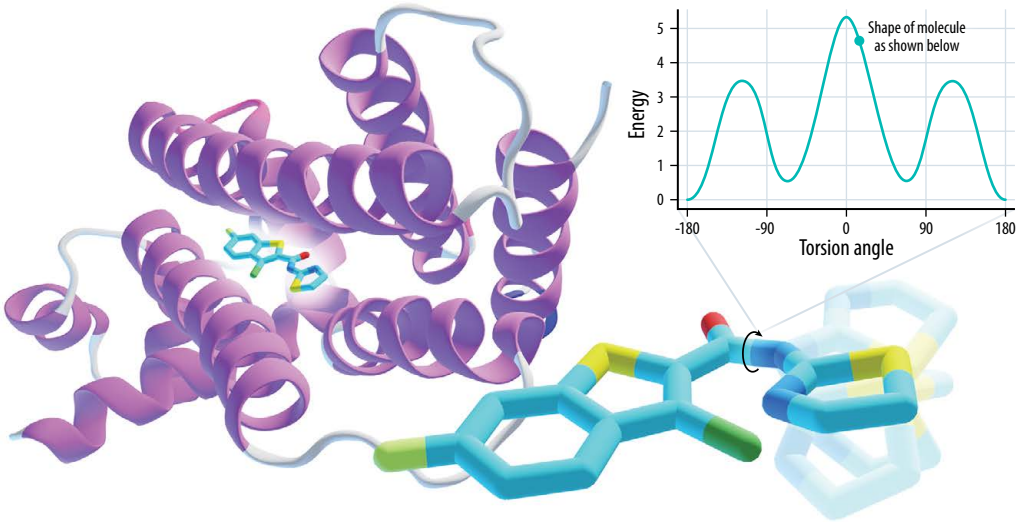

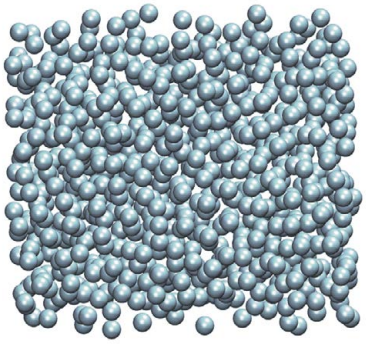

ML models can be repurposed through retraining. Here, the same method that was used for the pathogen-drug interaction simulation at left has been applied to study the properties of liquid aluminum. ML is better than classical computing methods at simulating disordered arrangements like the atoms or molecules of a liquid at a high temperature. *CREDIT: Justin Smith/LANL*

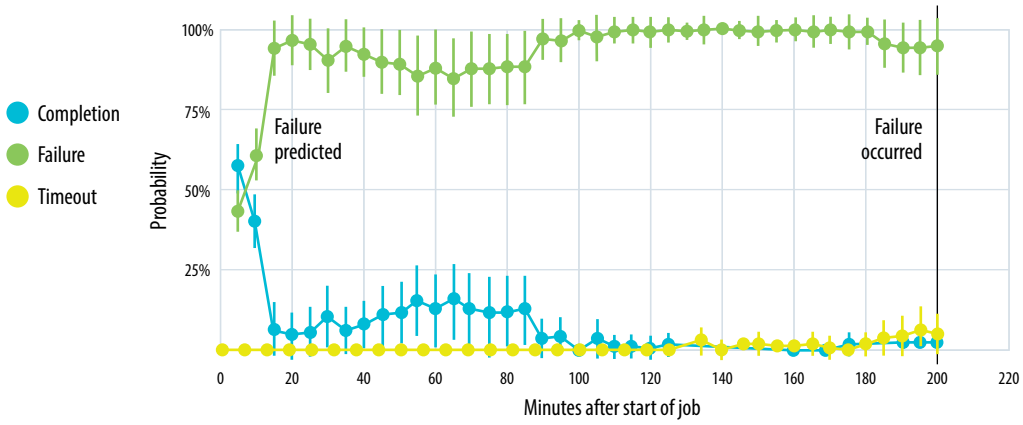

Lissa Moore is developing ML methods to improve high-performance computing by predicting job timeouts and system failures before they occur. In this proof-of-principle experiment, the ML tool was given a job that was known to have failed after about 200 minutes. Initially the model predicted equivalent likelihoods of completion or failure. By about 15 minutes in, nearly three hours before the job actually failed, the model began predicting a much higher probability of failure than success. This work has implications for improving the efficiency and reliability of supercomputing simulations.

trying to learn the names of a group of people, and each person wore the same color shirt many days in a row, you may inadvertently learn that "Bob" equals "red shirt" and "Barbara" equals "yellow shirt." But what happens when, one day, Bob and Barbara are both wearing green? This kind of false cue, called overfitting, can thwart ML as well, resulting in a model that has memorized rather than learned, and so is overspecialized and not transferable. The model has to incorporate known physics to answer known questions, but it also has to be transferable in order to get at unknown questions.

Whereas ML offers accelerated computation for atomistic systems, Lubbers' broader goal is to create models and methods that are transferable, so that they might be applied across many domains. Using transferable models can help scientists learn things they didn't know before and predict things they have never seen. It can help them define the very questions they need to ask.

The questions one has to ask are tied to the nature of knowledge.

"The questions one has to ask in scientific machine learning are tied to the nature of knowledge," Lubbers says. "If the algorithm learns, but the human doesn't, what has really been achieved? As we explore what's possible with machine learning, we are also learning how to approach a problem so that we will gain scientific knowledge as well."

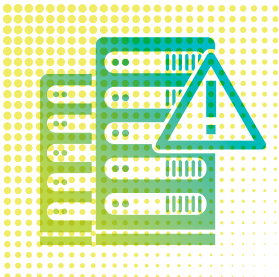

Better outcome

High-performance computing (HPC) simulations are one of the main ways that high-quality, high-stakes science gets done. And ML, whether physics-informed or not, can improve the reliability and performance of mission-central HPC simulations themselves.

"If the coin is the intersection of high-quality science with national security challenges," explains Laboratory ML expert and former Navy research scientist Ben Migliori, "then the two sides of that coin, the two ways we can use ML to that end, are, yes, embedding the physics in the model when we know the physics of the system, but also, when we don't actually know the physics and therefore can't embed it, understanding how and when it will fail."

Los Alamos ML expert Lissa Moore is developing methods to do just that. Not to be confused with out*put*, which is the answer to the question being calculated, the out*come* of a simulation is whether the simulation itself will run successfully, or if instead the system will crash, time out, or otherwise stall before completion. These kinds of interruptions increase the time and cost of HPC simulation, but preventing them requires knowing about the health of the HPC cluster and anticipating abnormal behavior.

An HPC cluster, or supercomputer, is basically 2000 computers networked together to do one thing. There are different ways to set it up, but no matter how it's set up, it needs continual monitoring. Both hardware—processors, networks, memory modules, etc.—and software—operating systems, user codes, job schedulers, etc.—get monitored, which generates terabytes of system-health data every day. Moore uses ML to take in all that data and make sense of it so that decisions can be made about if, when, and how a person should intervene to optimize the outcome.

Scientists queue for 12-hour time slots on the Laboratory's HPC clusters, so their codes have to run to completion in that time. If the operators who supervise the HPC clusters knew in advance that the system was going to crash, they could kill a simulation early and restart it, perhaps with enough time to complete a re-run. Or, if they knew that the simulation was going to time out, the operators could save the data midway, so that

when the user re-runs the simulation, it can pick up where it left off, rather than starting over and timing out again.

"This is important," says Moore, "because until now we've had only rudimentary tools looking for known problems. ML can look for new problems, ones we haven't anticipated or seen before. The machine learns what normal and healthy looks like, then tries to spot any deviation from the norm."

At present, the learning task is just recognizing potential problems. If the model predicts something unusual is going to happen, it raises an alert to a human who makes the call as to whether or not it is a real problem, and if so, what to do about it. The next step will be teaching the model to make that determination itself, and after that will come teaching it which actions to take in certain circumstances. Humans will still be in charge, but the ML tool will help them keep HPC simulation as efficient as possible.

Independently of her supercomputer-health monitoring work, but not unrelated, Moore also works on explainable machine learning. Firmly in third-wave AI territory, explainable ML is when the model not only makes a decision by itself, but can tell the user why it decided what it did.

What's going on under the hood of a trained ML model is typically not intuitive. Because the hidden layers are hidden and the algorithm is not defined by the programmer, the details of the mathematics in those layers are a bit of a black box. The programmer can learn about the trained model by sort of poking at it phenomenologically: vary the input and see what effect it has on the output. This is how one might discover that a model has learned Bob and Barbara's shirt colors, rather than their faces. But with explainable ML, the model could actually report back, "I know this is Bob because of the red shirt he's wearing," or, "I know this is Bob because I recognize his face." Being third wave, the field of explainable ML is still quite new and still rapidly advancing, but the potential is tremendous, offering a new level of performance and reliability for HPC simulations and other ML applications alike.

More like life

The Laboratory's Trinity supercomputer runs on 20 megawatts of electricity while the human brain runs on about one millionth of that, yet it's unclear which is more powerful. They can do roughly the same number of operations per second but they have very different skill sets.

"Here's the interesting thing about brains from an ML perspective," explains Los Alamos ML expert Andrew Sornborger, "Their main characteristic is that they learn. You send a brain out into the world, and it will interact with its environment and learn things without explicitly trying."

"But," adds Garrett Kenyon, a Laboratory physicist-turned-neuroscientist who collaborates with both Sornborger and Migliori, "no one really knows how brains do it. How we train machines is not how we ourselves learn."

Neuronal communication works through activity spikes. In a brain these spikes are ion-mediated charge fluctuations called action potentials, which send neurotransmitters across synapses to neighboring neurons. Spiking neurons offer more computing power for less electrical power, but most ML uses conventional, non-spiking communication schemes. Kenyon, Sornborger, and Migliori want to build ML models that operate more like brains, and they are all-in on spiking schemes. But spiking schemes don't map well to traditional computer architectures. So the scientists are using Loihi (low-EE-hee), a neuromorphic chip made by Intel, that directly emulates neural processing and synaptic plasticity in the brain by computing with biophysiologically-inspired spiking neurons.

Nearly all ML is supervised learning, meaning that a human tells the machine what its task is and when the task has been achieved. Brains also do supervised learning, but when they are brand new, they learn in an unsupervised way. In unsupervised learning, there is no specific task, other than taking in whatever data there is to take in and organizing it somehow. Just like toddlers learning to walk, human babies learning to see are not given stepwise instructions. As they are carried through the world with their eyes open and their brains on, they rapidly learn that things are different from other things and automatically create cognitive categories for all the things. The neurons of the baby's brain self-organize during unsupervised learning into a cortex that can process and represent all the incoming sensory data. This natural learning (unsupervised) is a prerequisite for the brain to be able to make use of the more structured training (supervised) the child will receive later, for example, when he or she goes to school.

A brain out in the world will learn things without explicitly trying.

Kenyon and Migliori are using Loihi as an unsupervised learning module that learns to represent visual input through self-organization. Just like a newborn, Loihi will be carried through the world seeing whatever there is to see (through a special camera), until it self-assembles all of its neurons into a kind of primary visual cortex.

Kenyon likens it to a digital crow. "Crows are very intelligent," he says. "A Loihi brain with a silicon retina, mounted on a drone, could be lightweight, low-power, solar-rechargeable, and fully autonomous. That's what I want to build."

But to assign a specific task to the digital crow, like "find blue cars and squawk when you see one," requires going back to supervised learning. By having done unsupervised learning first, the amount of training data needed for subsequent supervised training is drastically reduced. But, even though the amount of training data is reduced, the ML model still has to be "rewarded" and "punished" mathematically as it learns how to minimize the cost function. This can't be done on Loihi in the same way it is done on other processors, and that is where Sornborger comes in.

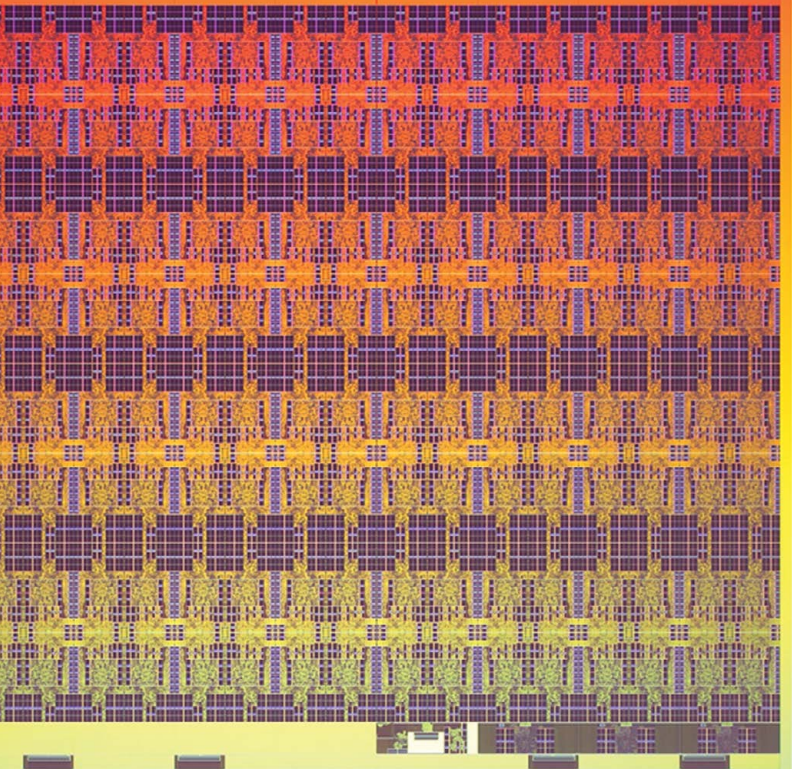

The Loihi chip, a neuromorphic processor introduced by Intel in 2017, is enabling the development of new machine learning models that emulate learning by human brains. The chip includes 130,000 neurons and 128 cores arranged in an architecture that is optimized for spiking neural network algorithms that mimic the electrical spiking behavior of neurons in human brains. *CREDIT: Intel*

On a DNN, doling out rewards and punishments is the feedback-motivated adjustment of weights and biases across the network, which is also called backpropagation of error, or simply backpropagation. But, until recently, backpropagation algorithms had not yet been implemented on Loihi, because of its unique, neurally based architecture. Brains have all the necessary information for changing their connections right there where the change occurs, at the synapse. Similarly, Loihi, as a brain emulator, implements local learning rules, where error is addressed right at the connection between neurons, rather than across the network. Sornborger and colleagues have taken a fundamentally nonlocal learning rule, backpropagation, and figured out how to implement it on Loihi, a system that only uses local rules.

How we train machines is not how we ourselves learn.

"This is what I'm most excited about," says Sornborger, "we've had backpropagation algorithms for years, but we didn't have the computers. Then we got the neuromorphic processor but couldn't do ML on it. Now that we've figured out backpropagation on Loihi, very big gains stand to be made."

These gains could be as pragmatic as drastically reducing the energy consumption of big server farms, as futuristic as the digital crow, or as lofty as a fundamental understanding of how our own human brains operate.

Making Connections

From solving old problems in new ways to finding new problems to solve, ML is a leading-edge technology that is only going to increase in popularity. This omnipresence is a boon to creative problem solving, but as different scientists across the Laboratory pursue ML to various ends, there is a risk of reinventing the wheel. To address this risk, a small team of researchers from three divisions across the Laboratory have been working for the past two years to coalesce what they've dubbed a "community of interest" in ML at Los Alamos.

"The community of interest consists of people who are developing ML as well as people who are interested in learning to use it," says Juston Moore, a Laboratory ML and cybersecurity researcher who first proposed the community-building effort. "It's difficult to implement ML in a safe and reliable manner, suitable for the Laboratory's critical national security mission. It should be a capability that is distributed across the Lab and a tool that is accessible to anyone who needs it. Domain experts should be able to vet their ML algorithms using a network of connected ML experts."

"We want to foster institutional knowledge," elaborates Migliori, who is part of the team, "just like the institutional physics knowledge we've amassed. Ideally we would have ML evangelists embedded in each group to help found ML projects."

Whereas the Laboratory's physics expertise has had 70 years to amass, ML is a relative newcomer, so Moore, Migliori, and their team want to speed things along.

There are hundreds of people across the Laboratory who work with or on ML, and the community-of-interest team has so far built two main mechanisms to help bring all those people together. First, they have launched a seminar series on adversarial machine learning—an emerging sub-field that uses competing attacker-defender models to strengthen performance—to get people into the same room and begin breeding familiarity. Second, they have established an internal topical chat service so people can converse in real time with ML colleagues, without clogging one another's email inboxes. Both efforts have been well received, word is getting around, and the network of ML people is solidifying.

There are many more projects that use ML at the Laboratory than the ones covered here. From improving cancer diagnosis to predicting earthquakes, from power-grid optimization to turbulence modeling, machine learning is revolutionizing how national security science gets done. LDRD

—Eleanor Hutterer

Earlier this year, a 20-nation team of scientists collaborating among 60 institutions, including Los Alamos, synchronized eight radio telescopes from around the world by atomic clock to make this first-ever image of a black hole. The team was recently awarded the Breakthrough Prize ("the Oscars of science") for the achievement.

This particular black hole resides at the center of galaxy M87. It is supermassive with the mass of about seven billion suns, and its gravity drives an extremely active surrounding region. As an enormous disk of gas spirals down toward the black hole, it acquires tremendous energy and radiates at various wavelengths, highlighting an otherwise invisible object. Yet both the complexity of this dynamic light source and the curved paths light takes in the vicinity of the black hole present a significant challenge for interpreting the observations. Los Alamos scientists Benjamin Ryan and George Wong contributed to this effort by modeling and performing supercomputer simulations (inset) of the inflowing gas and the way its glow would be warped by the curved space in order to generate a basis for understanding the actual black hole image. *CREDIT: Event Horizon Telescope collaboration*

ISSN: 1942-6631

Printed on recycled paper

1663
Mail Stop M711
Los Alamos National Laboratory
P.O. Box 1663
Los Alamos, NM 87545

1663magazine@lanl.gov

Presorted Standard
U.S. Postage Paid
Albuquerque, NM
Permit No. 532

A winter tableau typical of Northern New Mexico: This Pueblo Revival-style building in Santa Fe is adorned with *ristras* of dried chile peppers and festive *farolito* lanterns along the roof.

Los Alamos National Laboratory, an affirmative action/equal opportunity employer, is managed by Triad National Security, LLC, for the National Nuclear Security Administration of the U.S. Department of Energy under contract 89233218CNA000001.
A U.S. Department of Energy Laboratory LA-UR-19-32344

For a free subscription, visit:
www.lanl.gov/discover/publications/1663