Lawrence Livermore National Laboratory

March 2019

Science & Technology REVIEW

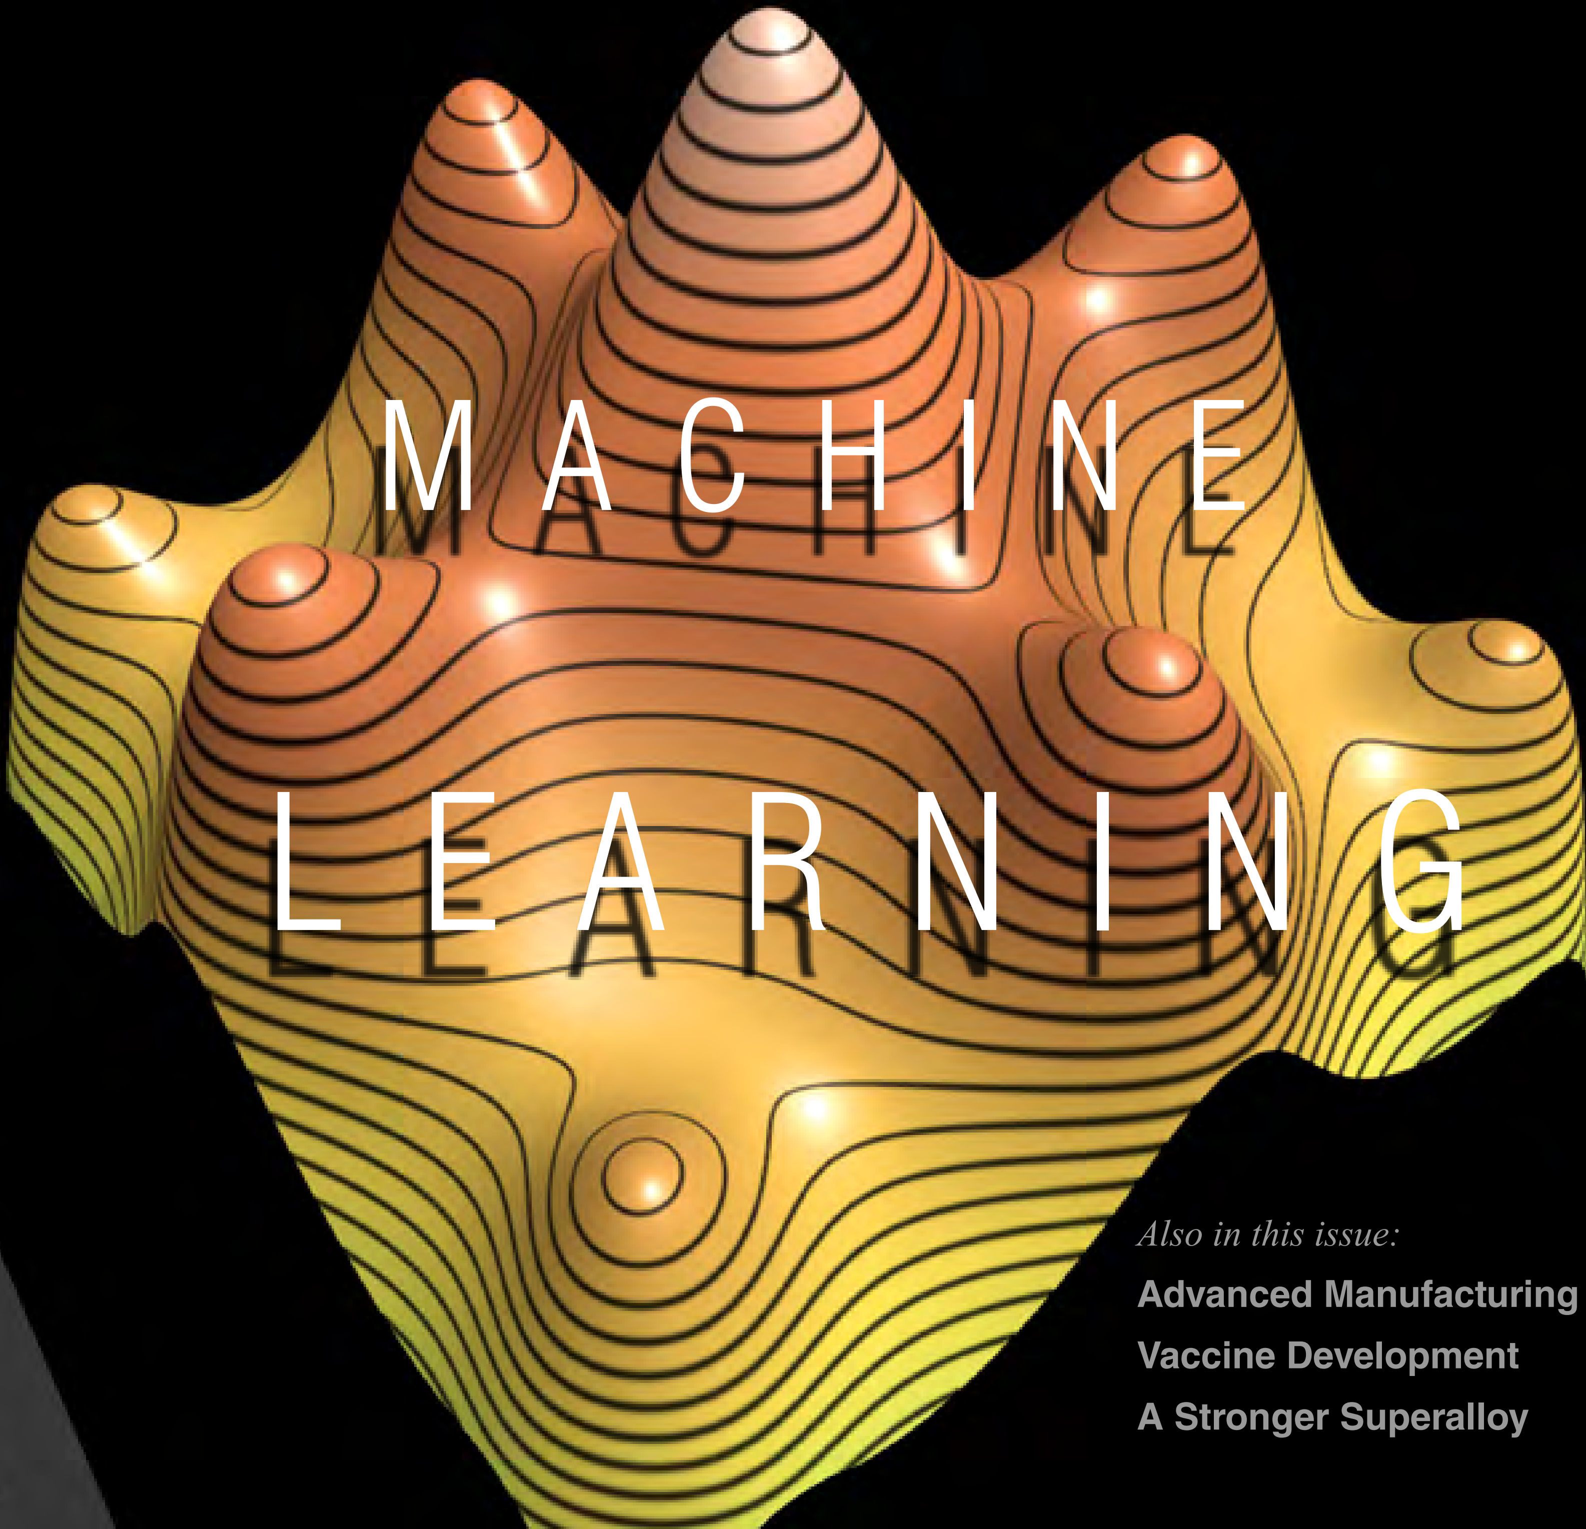

MACHINE LEARNING

Also in this issue:

Advanced Manufacturing

Vaccine Development

A Stronger Superalloy

About the Cover

Laboratory researchers are developing machine-learning (ML) tools to analyze vast amounts of data required for and generated by their mission-driven work. Key parameters in experiments and simulations are represented as separate dimensions in ML models. Topological data analysis is an important tool in understanding these high-dimensional parameter spaces. As the article beginning on p. 4 describes, these analyses help Livermore investigators explain and verify their ML models through visualizations that can reveal information about data that would not otherwise be evident. The front cover illustrates a mathematical function at the heart of an ML sample-optimization algorithm. The back cover shows visualizations from the National Ignition Campaign.

Cover design: Mark Gartland; cover art: Peter Lindstrom

About *S&TR*

At Lawrence Livermore National Laboratory, we focus on science and technology research to ensure our nation's security. We also apply that expertise to solve other important national problems in energy, bioscience, and the environment. *Science & Technology Review* is published eight times a year to communicate, to a broad audience, the Laboratory's scientific and technological accomplishments in fulfilling its primary missions. The publication's goal is to help readers understand these accomplishments and appreciate their value to the individual citizen, the nation, and the world.

The Laboratory is managed by Lawrence Livermore National Security, LLC (LLNS), for the National Nuclear Security Administration (NNSA), a semi-autonomous agency within the U.S. Department of Energy (DOE). LLNS is a limited liability company managed by Bechtel National, Inc.; the University of California; BWXT Government Group, Inc.; and AECOM Energy & Construction, Inc. Battelle Memorial Institute also participates in LLNS as a teaming subcontractor. Cutting-edge science is enhanced through the expertise of the University of California and its 10 campuses and LLNS' affiliation with the Texas A&M University system. More information about LLNS is available online at www.llnsllc.com.

Please address any correspondence (including name and address changes) to *S&TR*, Mail Stop L-664, Lawrence Livermore National Laboratory, P.O. Box 808, Livermore, California 94551, or telephone (925) 423-3893. Our e-mail address is str-mail@llnl.gov. *S&TR* is available on the Web at str.llnl.gov.

NNSA National Nuclear Security Administration

LAWRENCE LIVERMORE NATIONAL SECURITY LLC

Prepared by LLNL under contract DE-AC52-07NA27344

S&TR Staff

S&TR, a Director's Office publication, is produced by the Technical Information Department under the direction of the Office of Planning and Special Studies.

S&TR is available on the Web at str.llnl.gov

Printed in the United States of America

Available from
National Technical Information Service
U.S. Department of Commerce
5285 Port Royal Road
Springfield, Virginia 22161

UCRL-TR-52000-19-03
Distribution Category UC-99
March 2019

Science & Technology REVIEW

March 2019

Lawrence Livermore National Laboratory

Contents

Feature

Research Highlights

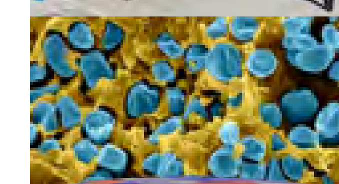

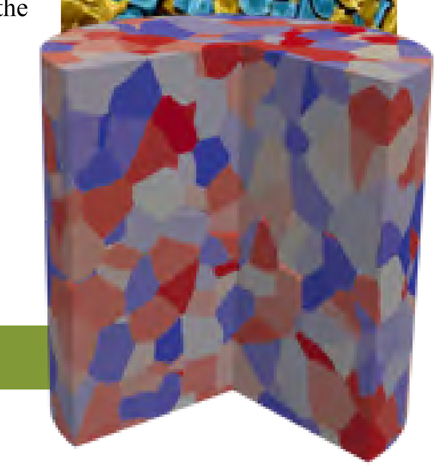

Departments

Petawatt Laser System Fully Integrated

The High-Repetition-Rate Advanced Petawatt Laser System (HAPLS) was recently declared to be fully operational at the European Union's Extreme Light Infrastructure (ELI) Beamlines. Forming the facility's third world-leading laser capability, HAPLS (see below) was developed and built by Lawrence Livermore to be the world's most powerful diode-pumped petawatt laser system. HAPLS has met the required performance parameters of being able to reach its 1-petawatt, 10-hertz design specification and is ready for integration with experimental systems.

In 2013, conceptual work on HAPLS commenced at the Laboratory, where it was designed, developed, and constructed by the Advanced Photon Technologies Program. In 2016, a Livermore–European team completed construction and final testing at the Laboratory. In 2017, the laser system was disassembled and shipped to ELI in Dolní Břežany, Czech Republic, where it arrived in June of that year.

The system consists of a main petawatt beamline capable of delivering 45 joules of energy per pulse, and is energized by diode-pumped lasers capable of delivering up to 200 joules of energy per pulse. The laser has been commissioned for early experiments at its phase 1 operation point of 16 joules and a 26-femtosecond pulse duration at a 3.3-hertz repetition rate, equivalent to a peak power of approximately 0.5 petawatts after the pulse compressor.

HAPLS represents a major advancement over any other petawatt-class laser system in the world and opens up a new arena of quantitative science, with the potential to develop societally impactful applications for industry and medicine.

Contact: Constantin Haefner (925) 422-1167 (haefner2@llnl.gov).

Understanding the Universe through Neutrinos

Livermore scientists are part of a team that proposed the next-generation Enriched Xenon Observatory (nEXO) experiment for providing a two-orders-of-magnitude increase over current limits in sensitivity to neutrinoless double-beta decay (NDBD) half-life. Determining features of the neutrino by observing NDBD—an extremely rare nuclear process—could provide an explanation for the puzzling overabundance of matter over antimatter in the universe. Studying NDBD could also reveal physics that would confirm the existence of a new elementary particle, the Majorana fermion. This discovery could reshape the Standard Model of particle physics and lead to a better understanding of neutrinos and their role in the universe's evolution. The design of the nEXO detector—a 5-ton liquid-xenon time-projection chamber using 90 percent enriched xenon-136—takes advantage of advanced technology for the next phase of NDBD research. The Laboratory research behind the experiment appears in the June 2018 issue of the journal *Physical Review C*.

"A competitive two-orders-of-magnitude increase in NDBD half-life sensitivity over current experiments is possible using the nEXO detector," states Livermore scientist Samuele Sangiorgio, lead author of the team's paper. "We now have great confidence in nEXO's design and approach, and we could have a real chance at measuring this rare event." Scientists expect to base discovery on observing only a dozen or so decays in a decade-long experiment. This very low signal rate means false signals from background radiation and cosmic rays must be suppressed as much as is feasible—a goal that the new experiment design will help achieve.

Contact: Samuele Sangiorgio (925) 422-6439 (sangiorgio1@llnl.gov).

Hydrodynamic Instabilities Elucidated

Researchers at Lawrence Livermore and the University of Michigan reported on recent experiments and techniques designed to improve the understanding and control of hydrodynamic instabilities in high-energy-density (HED) settings, such as those that occur in inertial confinement fusion (ICF) implosions at the National Ignition Facility (NIF). In the June 2018 issue of *Proceedings of the National Academy of Sciences*, the team describes four areas of HED research that focus on Rayleigh–Taylor (RT) instabilities. This phenomenon arises when two fluids or plasmas of different densities are accelerated together, with the lower density fluid pushing and accelerating the higher density one. RT instabilities can degrade NIF implosion performance by amplifying target defects and causing perturbations through engineering features, such as the tents used to suspend the target capsule in the hohlraum and the fill tube that injects fusion fuel into the capsule.

The paper summarizes a wide range of studies about HED RT instabilities relevant not only to ICF but also astrophysics, planetary science, and hypervelocity impact dynamics. The researchers state that the studies, although aimed primarily at improving understanding of stabilization mechanisms in RT growth within NIF implosions, also offer unique opportunities to study other phenomena that typically can be found only in high-energy astrophysics, astronomy, and planetary science, such as the interiors of planets and stars, the dynamics of planetary formation, supernovae, cosmic gamma-ray bursts and galactic mergers.

Contact: Bruce Remington (925) 423-2712 (remington2@llnl.gov).

Machine Learning and Laboratory Science Drive One Another

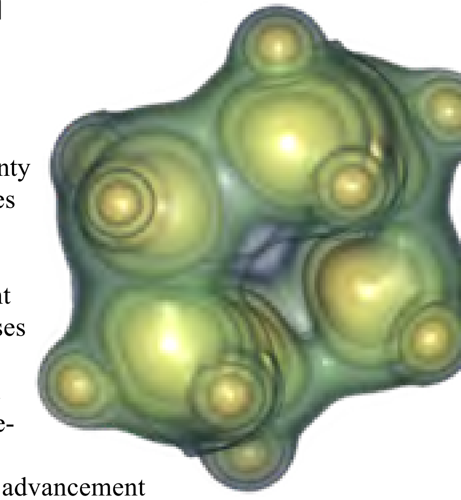

LAWRENCE Livermore is a data-rich environment. As the demand for sophisticated methods of analyzing and interpreting data grows, so too does the need to push the boundaries of data science—a field that goes beyond merely crunching numbers to exploiting sophisticated technologies such as machine learning (ML) to analyze data. Organizations all across the Laboratory apply data science techniques to scientific questions while also strengthening the field's methodologies. The article beginning on p. 4 describes this dual pursuit in ML, in which computer algorithms learn from data to identify patterns, make inferences, and predict outcomes.

ML is a rapidly growing specialty, particularly in nearby Silicon Valley, California, where consumer needs propel technological developments. Around the world, major ML-focused conferences receive thousands of paper submissions, and the pace of scientific publishing is staggering. At the Laboratory, we contribute to this progress because ML has important implications for scientific data analysis and for our national security missions. Although many companies have pioneered ML methods for commercial use, the Laboratory's problems often have different characteristics and demand unique solutions. Data-driven scientific decisions rely on researchers to create nuanced ML algorithms that can derive meaning from simulations, images, text, speech, video, and other data types. The article spotlights several ways in which Laboratory scientists adapt ML techniques to tackle application-specific challenges.

This adaptation can be incredibly difficult. For example, consider the ML task of classifying objects in an image, in which the computer tries to do what human eyes and brains can do easily. Indeed, the technology behind the ML algorithms that can successfully identify a cat, for instance, in a series of photographs is impressively complex. The computer must learn to distinguish feline-specific features from other pictured objects, the background, or interference. The process may entail pixel-level comparisons, color analysis, recognition of edges and boundaries, and more.

Now imagine programming a computer to detect the slightest three-dimensional anomalies in a two-dimensional scan of a vehicle at a border crossing. The computer must predict locations in the single image to search for unusual but nonspecifically defined objects that are not directly observable by humans. Compared to the cat-identification model, this cargo-scanning problem has fewer data to evaluate with more uncertainty and far greater consequences of error—such as incorrect or missing identification. Livermore is at the forefront of ML research that addresses real-world scenarios whose requirements go far beyond those satisfied by out-of-the-box ML tools.

Equally important is the advancement of ML as a science. Throughout history, many technologies have been developed and used before their mathematical and physical underpinnings had matured. Understanding the underlying principles of ML methods is important for reducing flaws, improving results, and enhancing our scientific knowledge. At Livermore, we strive to better comprehend how these methods work—such as how a specific prediction is made—so that we can have greater trust in them, and they can provide greater explanatory power. We seek assurance that ML models correctly extrapolate information and accurately reach conclusions. Unlike a movie recommendation, the decisions we make have significant consequences, and so we must have high confidence in them.

Fortunately, the Laboratory's culture of multidisciplinary teamwork helps maximize the potential of this new technology. Our computer scientists, data analysis experts, and domain scientists work alongside each other to a degree unmatched in industry or academia. This collaboration across traditional disciplinary barriers, combined with powerful supercomputing capabilities, enables us to drive and respond to evolving technologies such as ML.

No doubt the ML landscape will look very different over the next few years, and it will become ever more important to the Laboratory's missions in ways not yet foreseen. Given the increasing amounts of data being generated by experiments, simulations, and other sources, our researchers will continue to embrace and invent new data science and ML methods. This transformation of data-driven science is just beginning.

■ Bruce Hendrickson is associate director for Computation.

S&TR March 2019

MACHINE LEARNING ON A MISSION

Livermore computer scientists advance machine learning technology for scientific applications.

MORE than just a buzzword, machine learning (ML) has become part of everyday life. Social media platforms recognize faces in photos. Online stores recommend products related to shoppers' browsing and purchasing behavior. Smartphones offer word-completion suggestions based on users' texting habits. Search engines refine results after learning from users' past actions. Only with ML technology can self-driving cars adapt to moving traffic.

ML uses computers to learn from data and make predictions about the environment. As the world generates more data, interpretation becomes more difficult. Lawrence Livermore computer scientist Peer-Timo Bremer explains, "Humans reach a limit where they cannot perform the analysis anymore." A smart machine—one that adapts to new information on the fly—can speed up processing and analysis times and improve its accuracy in identification and prediction tasks. Although commercial and consumer applications of ML are numerous, Livermore's mission space also presents ample opportunities for exploiting ML tools, often requiring new development beyond standard applications. (See box p. 7.)

Indeed, Livermore faces unique challenges in advancing the ML arena. Bremer points out, "Commercial companies do not solve scientific problems, just as national laboratories do not optimize selections of movie reviews. We therefore build on commercial tools to create the techniques we need to analyze data from experiments, simulations, and other sources." ML algorithms must be scaled for high-performance computing (HPC) machines, and different types and varying volumes of data complicate matters. For example, one project may have access to thousands of patient health records, whereas another may only have data from a handful of National Ignition Facility (NIF) shots. Bremer continues, "A team may have to sort through genetic sequences, protein structures, energy spectra, x-ray images, or combinations of these." Other issues with scientific data include noise and imbalance—such as a handful of successful drugs versus millions of ineffective compounds—which will bias traditional data-driven models.

Along with Bremer, computer scientists Rushil Anirudh, Harsh Bhatia, Bhavya Kailkhura, Hyojin Kim, Shusen Liu, and Jayaraman Thiagarajan are go-to ML experts. They take a bidirectional approach, both advancing underlying theory and solving real-world problems. The algorithms involved are run on several on-site HPC resources, including Sierra, the Laboratory's newest and fastest supercomputer.

As valuable tools for analyzing data from scientific simulations and experiments, machine learning (ML) algorithms are run on many of Livermore's high-performance computing resources, such as the new Sierra supercomputer. (Photo by Randy Wong.)

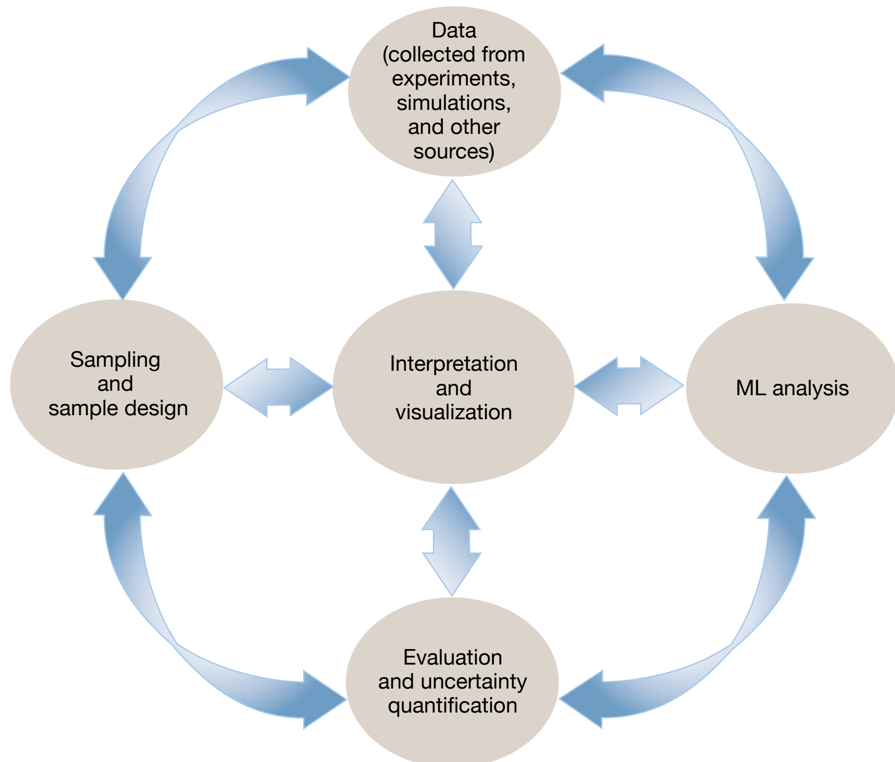

In scientific analysis using ML, sample design informs data collection in simulations and experiments. The ML model processes the data and generates predictions, which are then evaluated for quality. Results from both training and actual data are fed back into the sample design to refine the process. Visualizations are used for interpretation.

Perfecting the Process

Scientific analysis involving ML generally follows a cycle in which sample design guides data collection. Data are processed with ML algorithms and their associated frameworks—collectively, the ML model—which are designed to learn from data inputs. Results are scrutinized for errors and unknown variables, providing statistical quantification of uncertainties and informing subsequent sampling. All stages of the cycle are interpreted with visualization tools. The ML model is first trained on smaller, representative data sets to refine this process.

ML algorithms serve various purposes. For instance, neural networks (NNs) connect artificial neural units to observe and make inferences from data. Deep learning is another category of algorithms in which hierarchical layers of NNs adaptively learn from data to discover new features. In addition, ML methods respond differently depending on data properties. In supervised ML, the system analyzes labeled or classified data. In unsupervised ML, data are not labeled or classified, so the computer learns to identify common traits. Other types of learning are self-supervised—labeled and unlabeled data combined—and reinforcement—based on prior performance.

Livermore researchers actively develop new ways of configuring and deploying such algorithms. The common thread is improving ML's accuracy and efficiency for the benefit of the entire scientific analysis workflow. Accordingly, Thiagarajan explains, "All application domains face the same issue, and the conversation must start with the kind of data needed. Scientific analysis is driven by data."

Designing the Data Sample

Sample design is the key to quality results, especially when a project faces scope or resource constraints. Kailkhura notes, "Sampling requires configuring experiments and simulations to generate the most informative data. ML algorithms will not offer new insight if samples contain inaccurate or incomplete information." For instance, if a team with limited computing resources wants to run simulations of NIF shots, beforehand they must choose, as the focus of their investigation, the most valuable parameters. "These are the parameters," Kailkhura explains, "that will acquire the most information given a number of simulation runs, such as implosion dimensions."

Kailkhura looks at sample design abstractly, seeking mathematical solutions for sampling optimization problems. A high-dimensional (HD) parameter space is needed to represent the key factors that affect the results of a complex experiment or simulation. The higher the dimensionality, the more data are required to sample the space. Kailkhura describes these spaces from a theoretical perspective, citing the widely known sphere-packing problem—finding the ideal arrangement of oranges in a crate for *n* dimensions. In this problem, oranges represent data points in a sample, and the crate is the domain of interest, as in an inertial confinement fusion (ICF) implosion. The way the oranges are packed signifies the pattern of selected data points. Optimized sphere packing, or space filling, enables ML models to process data more quickly by minimizing the number of steps to reach a solution. Moreover, the models can provide insights into data not acquired yet, hence ML's predictive capabilities. Kailkhura seeks to cover as much of the space as possible while also obtaining the greatest information from the data sample. He states, "We strive for the right balance between coverage from uniform sampling and maximized information from random sampling. The optimal sample design will have some combination of uniformness and randomness."

Kailkhura collaborates with Bremer and Thiagarajan on a project, funded by the Laboratory Directed Research and Development Program, aimed at exploring spectral sampling of HD spaces. In this context, *spectral* refers to the frequency of change among data points—a necessary consideration, the team argues, in addition to the data's spatial arrangement. Spectral analysis can enable better understanding of space-filling sample designs by finding a balance between uniform and random coverage. The project's goals are to determine optimal sampling patterns and to create ML algorithms that can generate those samples in any HD space.

The project team uses a combination of exploration and exploitation techniques—sampling input variables independently of the output while using knowledge of the output to guide sample design, respectively. Project leader Thiagarajan notes, "We start the optimization process with blind exploration of data points, then switch to exploitation to search the regions of highest interest." This hybrid approach achieves better results than traditional methods by weighing both high- and low-frequency information—information respectively about more and less frequent change. The range of analyzable frequencies is maximized, providing statistically higher confidence in results. In recently published tests using data from NIF hot spot simulations, Livermore's spectral-sampling technique doubles the accuracy to significantly outperform other sample designs. Ultimately, optimized data inputs improve the ability of ML models to make useful predictions.

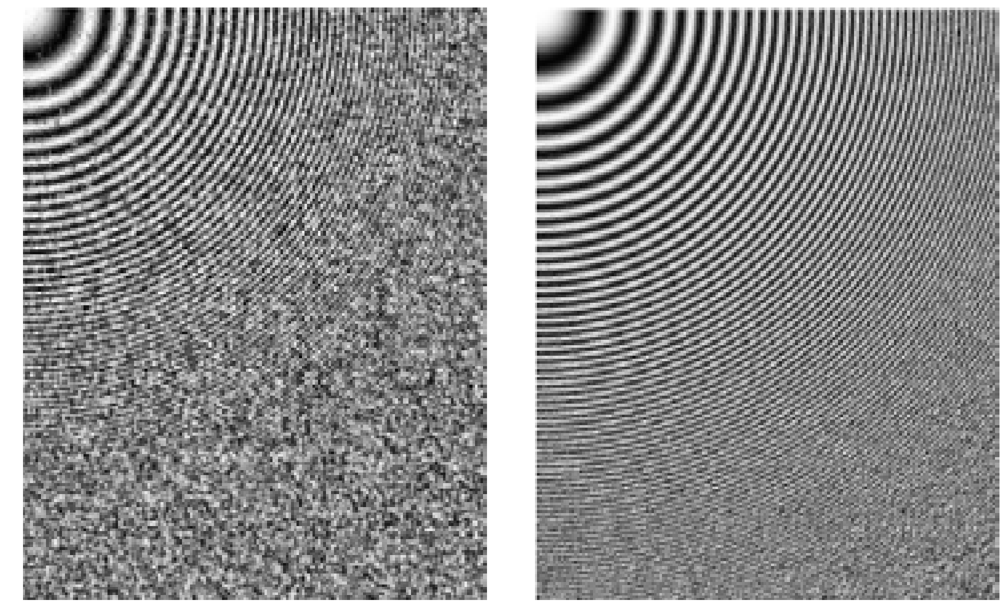

(left) Random sampling of image reconstruction data finds only one significant pattern, as shown by concentric rings (upper left corner). (right) The team's spectral-sampling method reduces noise and other artifacts to reveal additional zones of interest in the data.

Essential Expertise

Examples abound of Lawrence Livermore's growing demand for machine learning (ML) to solve challenges in scientific data analysis. One research team created a toolkit that trains massive neural networks on image data. (See *S&TR*, June 2016, pp. 16–19.) Another project focuses on time-varying data, which reveal patterns in time. In this scenario, new ML algorithms progressively use existing observation-based data to forecast future events. For example, clinical decision making could be enhanced by ML analysis of trends in patient data.

One National Ignition Facility project leverages ML to analyze the largest-ever data set from inertial confinement fusion (ICF) implosions. (See *S&TR*, September 2018, pp. 16–19.) Another group is developing an innovative cognitive computing platform that combines ML with graph analytics and other areas of artificial intelligence to improve ICF simulation efficiency.

ML is also speeding up data analysis and prediction in three-dimensional printing and making multimodal data analysis easier in nuclear nonproliferation. Materials scientists use ML and big data analytics to accelerate materials synthesis and optimization. (See *S&TR*, July/August 2017, pp. 16–19.) ML technology helps Livermore scientists catalog and interpret objects orbiting Earth and process huge volumes of data captured by ground- and space-based telescopes. Livermore has also partnered with several institutions to accelerate drug discovery and development by integrating high-performance computing, ML, and other data science technologies. Computer scientist Rushil Anirudh notes, "The possibilities of ML are exciting. Whenever we reach a roadblock, we find ways to break through with ML."

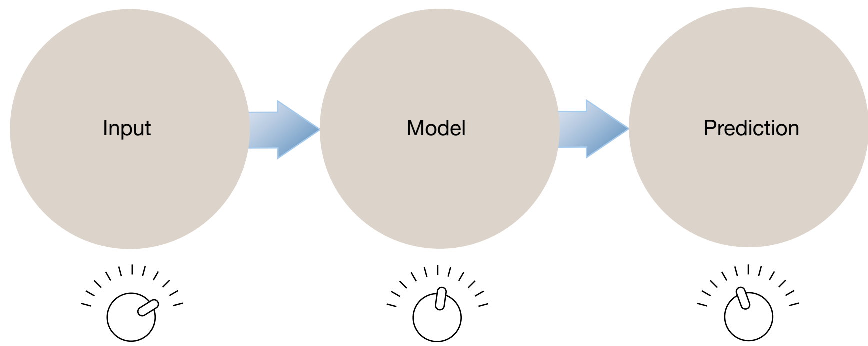

Perturbation is one method of exploring an ML model's interpretability in which researchers adjust—in a manner akin to turning a dial—different factors and observe the effects.

Trusting the Model

Model interpretability is another methodological ML pursuit at the Laboratory. "Human nature requires justification," states Liu. "We want to know which symptoms correlate to a particular diagnosis. We want to know how a conclusion was reached." Justification means providing rationale for how ML models work and the results they predict so stakeholders will trust both. Liu continues, "Interpretability is a necessary part of explaining or modifying an ML model, especially if the application is as important as NIF and not simply images of cats and dogs." Interpretability involves confronting tricky questions. For example, Bremer asks, "If a model is trained on a certain data set, how does one know it is not biased toward those data's properties?" An ML model might advise a bank that residents of a certain neighborhood are unsuitable candidates for a mortgage loan. If an applicant's address is the only criterion affecting loan approvals, then the model ignores other relevant information such as credit score or loan repayment history. Avoiding bias means understanding how the model arrives at a prediction and finding where bias might originate.

ML models do not have to specify a path to a solution. Consequently, Bremer cautions, "We may not understand how the model performed its analysis, which undermines confidence in the solution, especially for nonexperts." The stakes are even higher for large-scale models with thousands or millions of parameters. The quest for useful ML interpretation comes with many challenges, such as the absence of a universally agreed-upon explanation. To control error and variability in new ML approaches, Liu advocates for transparency so that the model is not merely a "black box."

Liu studies ML models through exploratory analysis. He states, "Conventional interpretation techniques study the model as an invariant object, where its behaviors are recorded and analyzed in an offline fashion." Instead, Liu recommends perturbation as one interpretation tool. Analogous to adjusting a radio's volume by turning a knob, researchers can perturb different variables and observe the behavior of others. For instance, masking a localized part of an image can affect the prediction of what the image contains. By dividing an image into a pixel grid and shifting the mask around the grid, researchers can calculate each pixel's importance in identifying the desired object. Liu explains, "This approach investigates relationships between inputs and outputs to determine which properties of the input contribute to the prediction."

Optimization of latent spaces presents another step toward interpretability. Latent space lies between ML processing's encoding and decoding stages and captures variations and other key underlying information in a compressed representation

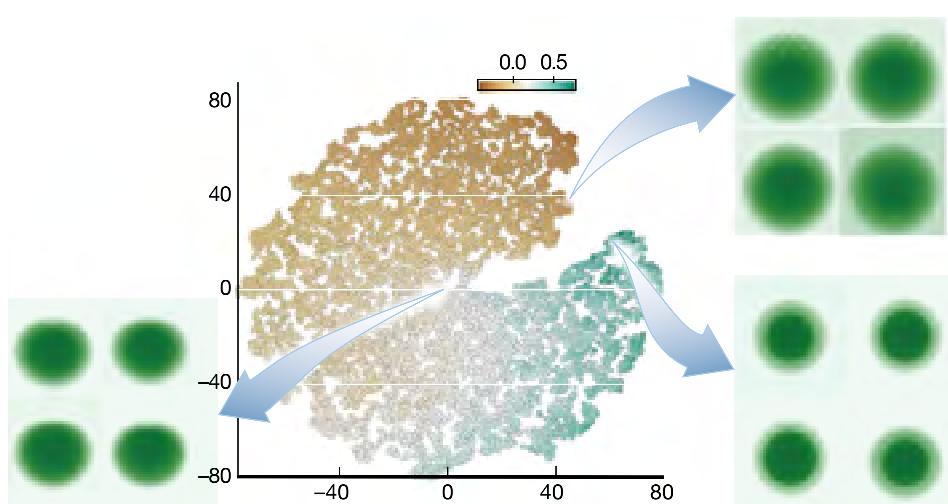

A 10-dimensional latent space of inertial confinement fusion (ICF) simulation data is reduced to the 2-dimensional visualization shown, in which the axes and scale no longer have explicit physical meaning. (insets) Different areas of the latent space capture various shapes of ICF images, providing insight into how the ML model interprets variations in high-dimensional (HD) data.

of the data. Unsupervised ML models map inputs through layers of NNs into the latent space, where data are reduced into lower dimensional representations, enabling the model to identify hidden features beyond those observed. "In many real-world scenarios, HD data can be compressed into spaces with as few as two to four dimensions," notes Kailkhura, whose work in sample design optimizes an understanding of these latent spaces. Knowing more about the features of these latent spaces makes results more interpretable.

"Latent spaces are compact and descriptive but typically not transparent or intuitive," says Bremer. Therefore, Liu and colleagues apply nonlinear dimensionality reduction functions to latent spaces and use visualizations to discover feature variations captured and distributed throughout these spaces. By comparing visual encodings of the HD space, researchers can determine how many dimensions yield the most valuable information. In one study of ICF simulations, the team compared image patterns in 10- and 16-dimensional latent spaces and found that the latter did not fully use all dimensions. Liu summarizes, "By reducing the dimensions when exploring the latent space, we can directly assess the information captured by that space and explain the differences between simulations."

For many scientists who rely on ML, seeing is believing. Topological data analysis is another valuable tool for understanding the structure of HD spaces, and the resulting visualizations help Livermore researchers explain and verify ML models. "Topology produces abstract structures that generalize to high dimensions," notes Bremer. Laboratory researchers have released open-source software that render data relationships through mountains, valleys, and other maplike contours. Bremer continues, "We can extract HD properties and show them as a low-dimensional terrain. Visualizations allow us to find patterns or anomalies that other statistical methods may not find, so we can evaluate information that would otherwise be incomprehensible."

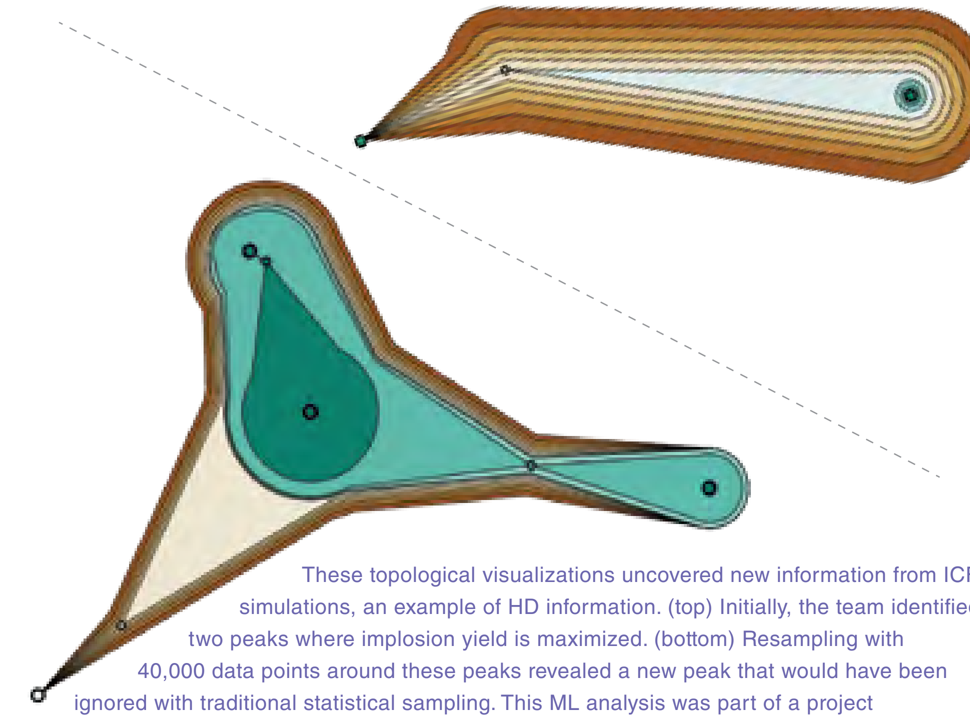

These topological visualizations uncovered new information from ICF simulations, an example of HD information. (top) Initially, the team identified two peaks where implosion yield is maximized. (bottom) Resampling with 40,000 data points around these peaks revealed a new peak that would have been ignored with traditional statistical sampling. This ML analysis was part of a project investigating ICF target shape. (See *S&TR*, September 2018, pp. 16–19.)

Case Study: Multiscale Modeling

In 2016, the Department of Energy (DOE) and the National Cancer Institute launched a multiyear partnership to advance cancer research using modern HPC resources. Livermore plays a central role in the program's three pilot projects. (See *S&TR*, October/November 2016, pp. 4–11.) One project, nicknamed Pilot 2, brings together three DOE laboratories—Lawrence Livermore, Los Alamos, and Oak Ridge—and Frederick National Laboratory for Cancer Research to explain interactions between cell membranes and specific proteins that induce many forms of cancer. (Pilots 1 and 3 focus on drug discovery and patient health records.)

To guide multiscale simulations of these interactions, Pilot 2 collaborators—including Bhatia and Bremer—develop ML approaches aiming to understand both the mechanism of a protein called RAS and the signaling chain that causes another protein, RAF, to interact with RAS. "ML is at the very center of this project, integrating different areas of expertise," states Bhatia. "We use ML to locate phenomena occurring on the cell membrane in coarse simulations, which we can then investigate more closely with higher fidelity simulations." With this computational steering approach, researchers guide simulations to gain specific insight while maximizing the throughput of computational resources.

Understanding protein biology requires modeling at different spatial and temporal scales—from nano- to milliseconds and from nano- to micrometers. Bhatia explains, "Simulating the underlying phenomena with sufficient accuracy at fine scales is prohibitively expensive computationally." Therefore, the project's sophisticated ML model is trained on coarse macroscale simulations before resources are spent on more detailed microscale molecular dynamics (MD) simulations. He continues, "Coarse simulations give us a reasonable approximation of results. The ML model identifies important areas, such as a small location where a protein interacts with the cell membrane, where we should invest our resources at a higher resolution. We want to investigate enough regions of interest to make statistical claims over a long temporal range without running simulations for the entire period.” This tactic could cut simulation time from months to days.

Computational steering is a sampling problem. Accordingly, the approach takes advantage of latent spaces. Bhatia says, “We have millions of potentially interesting data points with nonlinear, highly complex relationships. For example, consider finding similar-looking houses among millions of photos. We could not simply compare individual pixels to determine similarity.” The team’s ML solution includes an autoencoder—a deep NN—that reduces the data into a latent space. From there, the model chooses the features most dissimilar from previous iterations and ranks the results according to importance—from most to least anomalous. Even with compression, a million data points could be flagged, which is why using latent spaces is key.

In a process called adaptive sampling, data generated by macroscale simulations inform sampling of the MD simulations—the latter, in turn, becomes part of the feedback loop to update the former. Together, autoencoding, adaptive sampling, and the in situ feedback cycle allow the team to manage over a million samples through HD analysis and, therefore, run macro simulations with the accuracy of MD. “These types of simulations are novel, and we are scaling the workflow to target a supercomputer such as Sierra,” states Bhatia. In 2018, the Pilot 2 team reached a major milestone by computationally steering such multiscale simulations on Sierra.

Case Study: Threat Detection

Three screening scenarios—medical diagnosis, airport luggage, and commercial truck cargo—share important characteristics. All require expert analysis of image scans, and threats are reduced with quick identification of suspicious objects. Automatically flagging suspicious areas in an image saves human operators’ time while minimizing errors. For example, maximized information from a computed tomography scan can help reduce a patient’s radiation exposure or improve prognosis with early cancer detection. In all three scenarios, the goal is higher detection rates with fewer false alarms.

Lung cancer nodules are inconsistent in size and shape and may not appear clearly in a lung image. A radiologist can mark a nodule in an image but cannot be expected to label every affected pixel. ML algorithms require more specific coordinates for nodule location, so the model must learn to create detailed labels at different stages of analysis. Anirudh explains, “We use unsupervised strategies to estimate nodule characteristics, such as boundaries, in weakly labeled data.” Kim adds, “This model will not replace radiologists’ expertise but will significantly reduce their workloads by filtering out images that do not need close review. The model can also provide a ‘second opinion’ to reduce diagnostic errors.”

Luggage screening technology stands to benefit from ML-driven efficiencies in image quality, as airport scanners often provide sparse views of imaged objects. The Livermore team built a system of one- and two-dimensional NNs to recover limited-angle or partial-view images. Mindful of interpretability, they also designed a confidence score to gauge results reliability. The score is calculated from estimates of pixel variabilities within the model’s latent space and is correlated with reconstruction quality. The team’s image reconstruction and segmentation techniques have shown higher fidelity to ground truth than other methods.

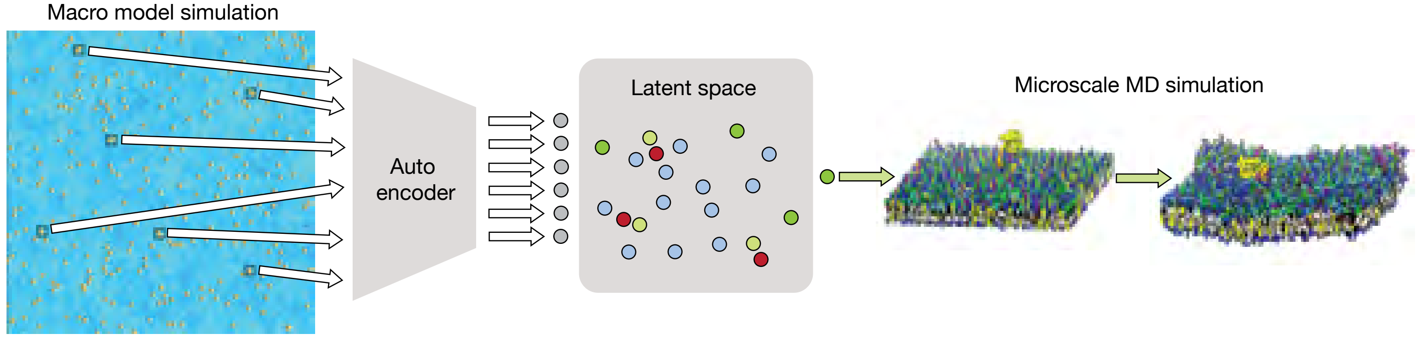

An illustration shows how ML is at the center of the Pilot 2 cancer research workflow, connecting coarse macroscale simulations (left) and fine microscale molecular dynamics (MD) simulations (right) used to investigate mediation of cancer initiation by the protein RAS. An autoencoder reduces macroscale model data into a latent space, where the data are ranked by novelty and importance. MD simulations are reserved for the most important regions, thereby conserving computational resources. The result is simulations with macro (long) length and timescales that also provide insights at the microscale.

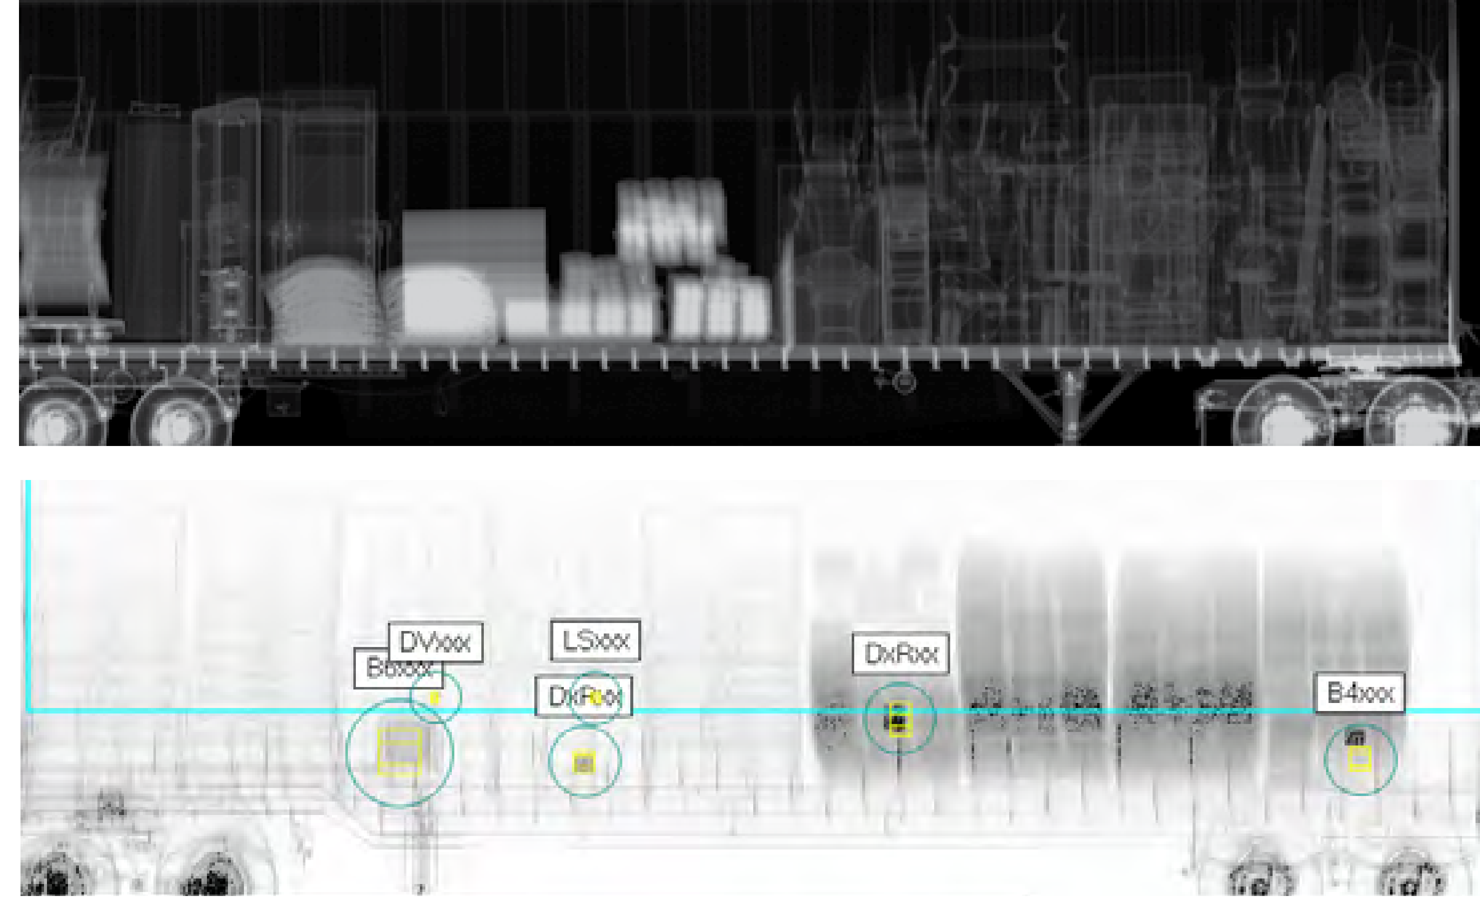

(top) Illicit materials inside a vehicle are difficult to discern from this single-view scan. (bottom) The Laboratory’s ML source-separation technique divides the image into layers from which the three-dimensional depth of the contents can be discovered.

At ports of entry, analysts see only a two-dimensional scan and must decide whether cargo contains, for example, nuclear materials stashed among a truckload of appliances. The cargo’s three-dimensional depth cannot be directly observed and so is inferred from a sum of the layers in a single image. The Livermore team has developed a source-separation model that splits a single image into multiple images to predict distribution of cargo materials. By training on probabilistic “clean” data separated into layers, this unsupervised ML model develops a surrogate for physical materials, then applies it to subsequent scans. Thiagarajan compares this technique to the way the brain identifies merged objects, saying, “If I show you separate images of a face and a pair of sunglasses, you can mentally combine them.”

In addition, Livermore researchers are moving toward what Anirudh calls “ML 2.0”—a more robust unsupervised model that does not collect data sets for every task. For example, Kim explains, “Thousands of unlabeled bags are scanned daily at airports. When a new object is introduced, the scanner needs to detect the abnormality for the security officer to investigate.” The team is solving such inverse problems with adversarial NNs, which perform generative and discriminative evaluations to reduce errors.

Disruptive Advances

In the quest to understand human intelligence, researchers across the Laboratory are evolving scientific ML in areas such as mathematical neuroscience, brain-inspired network architectures, representation learning, and multistage training algorithms. Thiagarajan says, “Combining scientific exploration and artificial intelligence opens up exciting opportunities for solving real-world challenges.”

Livermore’s ML experts agree that most research teams at the Laboratory will eventually seek ML-driven solutions to the challenges they face. In fact, many mission-critical programs already rely on ML technologies. Kailkhura states, “Once you grasp the concepts, the applications are numerous.” Bremer adds, “Grand challenges in science and computing cannot be addressed with incremental improvements. Instead, we must look for disruptive advances with significant technical, programmatic, and strategic impact. Livermore is absolutely the right place—perhaps one of the only places—to do this.”

—Holly Auten

Key Words: algorithm, computational steering, deep learning, high-dimensional (HD) space, high-performance computing (HPC), image reconstruction, image segmentation, inertial confinement fusion (ICF), Laboratory Directed Research and Development Program, latent space, machine learning (ML), model interpretability, neural network (NN), sample design, simulation, source separation, spectral sampling, topological visualization.

For further information contact Peer-Timo Bremer (925) 422-7365 (bremer5@llnl.gov).

S&TR March 2019

The Laboratory, industry, and academia are working side by side to develop advanced materials and manufacturing processes at the Advanced Manufacturing Laboratory (AML), located on the Livermore Valley Open Campus. Partnerships at AML are designed to further Livermore's missions and benefit partners developing commercial products and processes. (Photo by Randy Wong.)

THE ADVANCED MANUFACTURING LABORATORY IS OPEN FOR BUSINESS

ADVANCED technology is reshaping and transforming manufacturing the world over. Signs of this transformation are everywhere: factory automation, machine learning, additive manufacturing, robotics, and cloud-based process management, to name only a few trends. Livermore contributes to this renaissance through its research and development (R&D) in fields such as advanced manufacturing, partnerships that use the Laboratory's high-performance computing to improve industrial processes, and commercializing new manufacturing technologies.

Now, the Laboratory has opened the Advanced Manufacturing Laboratory (AML), a 1,300-square-meter facility where Livermore scientists and engineers are working side by side with partners in the private sector and academia to create new materials and technologies. R&D at AML aims to further Livermore's national security missions while enabling partners to release new products and services into the marketplace, a process called spin-in/spin-out technology development.

"U.S. industry is becoming highly innovative in manufacturing technology," says Patrick Dempsey, director of strategic partnerships in the Laboratory's Engineering Directorate. "AML will help us take advantage of industry's innovation by partnering our R&D efforts with theirs, accelerating the progress of both." Livermore's growing research efforts in additive manufacturing, along with development of an open space for research collaboration, planted the seeds that germinated into the new facility. From planning to opening, creating AML took approximately three years.

Outside the Fence for Easier Access

AML is located "outside the fence"—on the Livermore Valley Open Campus (LVOC), which is sited beyond the security fence. The LVOC location facilitates collaboration and communication by freeing partner personnel from the strict security requirements that must be followed when working within Laboratory boundaries. At the same time, AML's partnership model provides mechanisms to address the partners' concerns, including intellectual property and confidentiality.

"We are looking for partnerships with industry in areas of research and development that we both care about," states Chris Spadaccini, director of the Center for Engineered Materials and Manufacturing and head of Livermore's Additive Manufacturing Initiative. "We expect that both sides will contribute actively to efforts that advance technologies useful to our partners' needs and the Laboratory's missions. We are also pursuing partnerships with academic institutions, because AML could serve as a hub to support student and faculty interactions with Livermore and the commercial sector." In developing AML, the Laboratory took active steps to fully understand industry's needs, such as holding an "industrial partnerships day" to hear from potential partners in person and using the Federal Business Opportunities channel to solicit partnership ideas from companies. Anantha Krishnan, the Laboratory's associate director for Engineering, says, "AML is enabling a new business model for Livermore to team with industrial and academic partners, and the success of these collaborations could provide a template for public–private partnerships in the future."

Five Tracks to Cover the Possibilities

The result of this consultative development is a flexible facility with a strategic plan that offers five tracks for partnerships: design, materials, processes, applications, and qualification and certification. Each track focuses on a particular area of manufacturing relevant to a party's interests. In the design track, for example, partners can leverage Livermore's design optimization capabilities and high-performance computing resources to improve complex, multifunctional designs for their products. The materials track applies Livermore's expertise in developing new manufacturing materials, such as metallic particles, nanomaterials, glass, and liquid photo resins.

Lawrence Livermore National Laboratory

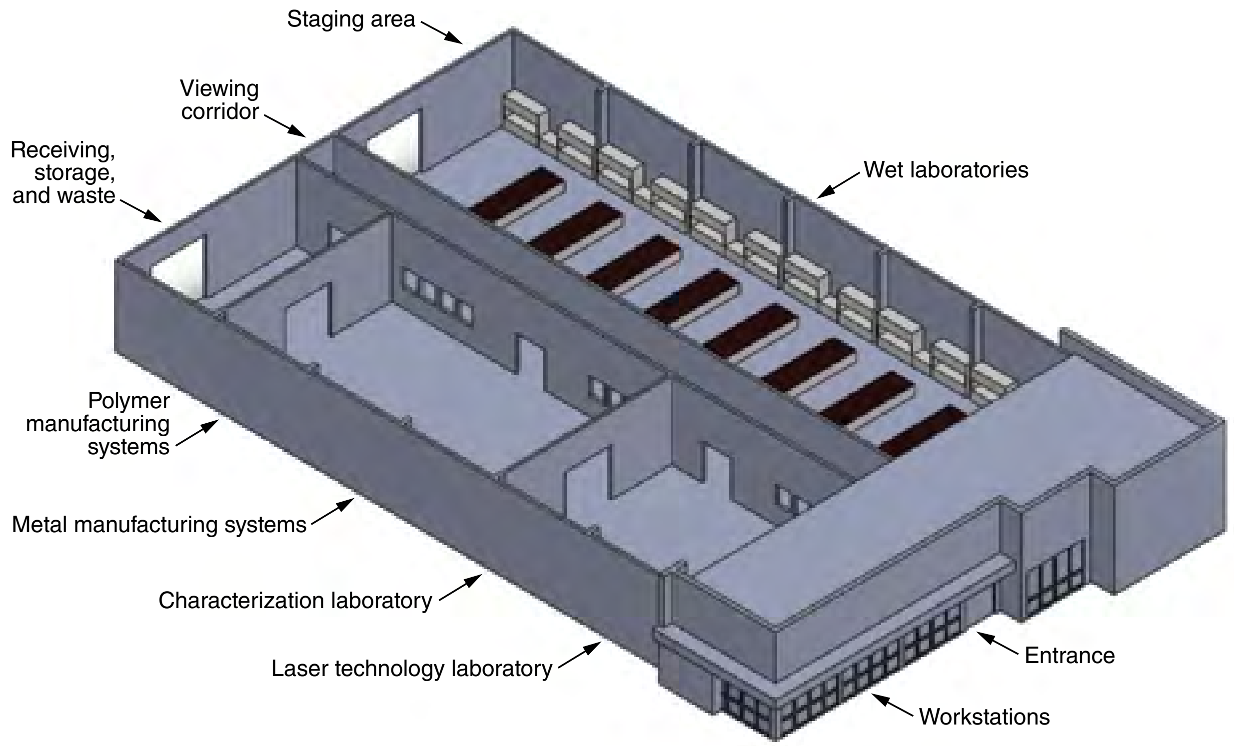

The 1,300-square-meter AML has some of the world's most sophisticated capabilities for developing new advanced manufacturing technologies, including wet and dry laboratories, laser-based technologies, and a characterization laboratory.

AML's facilities reflect Livermore's process expertise across a broad spectrum of materials and scales, including direct ink writing, powder bed fusion, electrophoretic deposition, projection microstereolithography, and laser-based processes such as two-photon lithography and selective laser melting. With these capabilities, partners on the applications track can develop new materials and components for nearly any sector—transportation, defense, energy, or biomedicine, for example. With some advanced manufacturing processes promising shorter development times, the qualification and certification track is tailored to accelerate the commercial acceptance of new materials, processes, and components. AML is currently home to more than half a dozen partnerships, and more are in the pipeline. Three examples demonstrate how the private sector is taking advantage of AML's full potential.

Locating Accurately in Time and Space

The Global Positioning System's (GPS's) network of satellites has revolutionized navigation, allowing even ordinary consumers with a GPS-enabled device to accurately determine their location almost anywhere on Earth. However, in a critical application such as aeronautical or marine navigation, failure of a GPS receiver could be disastrous. Vector Atomic, a startup in Oakland, California, is developing quantum sensors for inertial navigation and timing that do not rely on GPS. Instead, the devices will determine the position and orientation of a vehicle by precisely measuring its linear and rotational acceleration through the quantum interference of atomic wave functions. To achieve small size and low power usage in these devices, Vector Atomic is using AML's facilities and staff to incorporate Livermore-developed micromirror array technology, which controls and directs light using microscale structures to manipulate arrays of tiny mirrors. (See *S&TR*, September 2017, pp. 16–19.) In addition, the partnership will rely on AML's unique micro- and nano-additive manufacturing processes and precision microassembly capabilities.

Livermore is pursuing the Vector Atomic partnership through a cooperative research and development agreement (CRADA). Engineer Robert Panas, the Livermore lead for the partnership, explains, "The CRADA will allow us to merge these two technologies. Our micromirror technology can modulate and tune the laser power for the atomic sensors more accurately than any other available technology. The result will be a device that can operate in demanding environments with substantial jostling and shaking. We expect the combined product to represent an improvement of more than an order of magnitude over current inertial navigation systems."

Optics Polished to a Fine Shine

Large fluxes of energy pass through the optics of the 192 beams comprising the world's highest energy laser, Livermore's National Ignition Facility (NIF). With repeated shots, tiny pits in NIF's optical glass can enlarge to become damage sites that compromise performance. "The Laboratory has made a substantial effort to mitigate damage precursors and initiated damage sites on NIF's large optics," says Ibo Matthews, a group leader in Livermore's Materials Science Division. "The damage mitigation process we developed uses carbon dioxide lasers to repair damage on the surfaces of silica optics, smoothing their imperfections. We realized that this process could be used for the laser polishing of glass, even the localized repair of NIF optics." The enabling research was funded primarily by the Laboratory Directed Research and Development Program. (See *S&TR*, April/May 2017, pp. 17–20.)

At a conference in 2014, a presentation about the technology by Matthews, Materials Science Division staff scientist Nan Shen, and their team attracted the attention of Edmund Optics, which quickly entered into talks with Livermore. The U.S.-based company eventually established a CRADA to work with the Laboratory at AML. The partnership's goal is to extend Livermore's technology into a commercial system capable of polishing industrial lenses and mirrors to the same high surface quality demanded by NIF.

Perfecting a Manufacturing Control Method

An industrial manufacturing process must be carefully controlled to ensure that the final product conforms exactly to specifications. Livermore researchers have teamed with General Electric (GE) to explore a method called feedforward control in the additive manufacturing of three-dimensional (3D) parts. "In feedforward control, a computer model simulates a part and its manufacturing process," explains Wayne King, the partnership's Livermore lead and overall lead of the Accelerated Certification of Additively Manufactured Metals Project. "The simulation actually trains the manufacturing tool in building the part, so the quality of the final product depends on the fidelity of the simulation."

In each area of AML, advanced manufacturing equipment is being moved in and set up. Here, Laboratory researchers will collaborate closely with partners in industry and academia in nearly any sector, including defense, transportation, energy, and biomedicine. (Photo by Randy Wong.)

At AML, the partners are developing the prototype for a feedforward control–based machine capable of 3D printing high-quality metal parts. "We aim to get feedforward control into an additive manufacturing system within two years. The key to success is incorporating this control approach into a commercial machine, something that would not be possible without GE," says King. (See *S&TR*, January 2015, pp. 13–18.) The Department of Energy's Technology Commercialization Fund is supporting the Livermore–GE partnership. In addition, the Laboratory, GE Global Research, and several other partners are developing feedforward methods to speed up the qualification and certification of 3D-printed metal replacement parts for the U.S. Navy, an effort funded by the Office of Naval Research.

These three examples represent the many partnerships that are active or under development at AML—and the Laboratory is seeking to establish more. When Livermore's talented researchers partner with their industrial and academic counterparts in a space containing some of the world's most advanced manufacturing capabilities, the expected result is not merely accelerated innovation, but innovation that benefits U.S. industry and advances the Laboratory's missions in the face of rapid transformation in these arenas.

—Allan Chen

Key Words: additive manufacturing, advanced manufacturing, cooperative research and development agreement (CRADA), Edmund Optics, feedforward control, General Electric (GE), inertial navigation, Laboratory Directed Research and Development Program, Livermore Valley Open Campus (LVOC), micromirrors, quantum sensors, two-photon lithography, Vector Atomic.

For further information contact Patrick Dempsey (925) 423-1868 (dempsey4@llnl.gov).

S&TR March 2019

A LONG SHOT PAYS OFF IN VACCINE DEVELOPMENT

Lawrence Livermore biologists (from left) Amy Rasley, Nicholas Fischer, and Matthew Coleman pose in front of a flow cytometer. Using this system, the team performs multicolor flow cytometry and high-throughput cell sorting to interrogate immune cell populations of interest. Their aim is to understand the mechanisms by which their vaccine formulation protects against *Francisella tularensis*, the bacterium that causes tularemia. (Photo by Randy Wong.) (right) *F. tularensis* is highlighted blue.

VACCINES teach the immune system to fight disease by mimicking what the body would encounter during a natural infection. The immune system's response to a live bacteria or virus is similar to its reaction to the attenuated (weakened) or inactivated (dead) version of a pathogen used in a vaccine. After vaccination, reexposure to the live microbe causes the body's immune system to "recall" the prior reaction, stopping the infection more quickly than it would have without the initial encounter. This immunological memory is the basis for vaccine-mediated protection.

Many licensed vaccines use attenuated viruses or bacteria, which stimulate a strong response and confer robust immunity, but these types of vaccines have some drawbacks. Individuals with immune deficiencies may be unable to receive them safely. Furthermore, the attenuated microbe, once inside the human body, could mutate into a more virulent form—although this scenario is unlikely. In addition, live vaccines must be stored at low temperatures and therefore may be impractical for areas with limited refrigeration, such as war zones and developing regions. Inactivated vaccines, on the other hand, cannot mutate, are safe for nearly everyone, and can often be transported without refrigeration. Unfortunately, most inactivated vaccines, which provoke a weaker immune response than do live vaccines, may require booster shots to maintain immunity and so may provide only partial protection against exposure.

A different type of vaccine—one that uses components of a pathogen rather than the whole organism—could offer the best

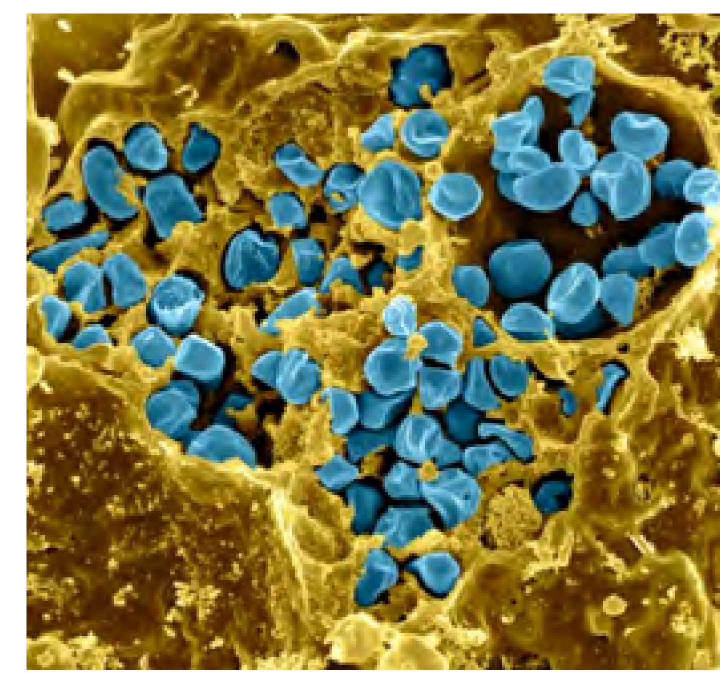

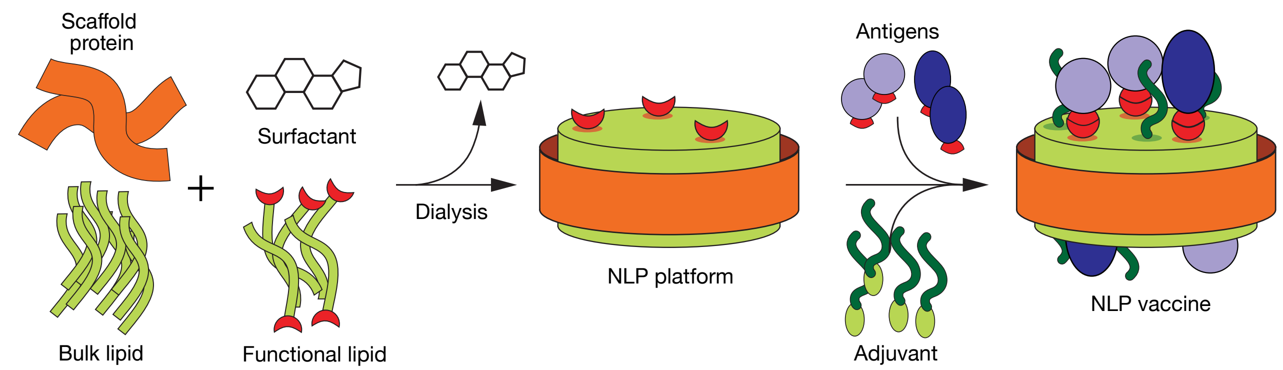

Nanolipoprotein particles (NLPs) are produced by mixing scaffold proteins and lipids with a surfactant. When the surfactant is removed through dialysis, the proteins and lipids self-assemble into a disklike structure. The resulting NLP provides a platform for attaching other molecules to produce a potent, safe, and targeted vaccine.

of both worlds. Subunit vaccines incorporate only those parts of a disease-causing microbe that best activate the body's immune system against the pathogen—namely, uniquely identifiable marker molecules called antigens. The potential advantages of such a vaccine are substantial—less variability between doses, fewer adverse or unintended reactions, and more freedom to formulate the vaccine for optimal immune response. Of course, subunit vaccines present their own challenges for researchers, particularly in identifying the right antigen formula to protect against a pathogen and finding the most efficient and effective method of delivery.

As part of efforts to defend soldiers and first responders against bioterrorism and biowarfare, Lawrence Livermore researchers have made exciting breakthroughs in subunit vaccine research. In vivo rodent studies, for instance, have shown that Livermore's nanoparticle delivery platform increases the efficacy of subunit vaccines, making the approach viable for a range of bacterial and viral threats. More recently, with support from the Laboratory Directed Research and Development Program and the Defense Threat Reduction Agency, a team of Laboratory scientists and their academic collaborators became the first to demonstrate a subunit vaccine capable of providing 100 percent protection in test animals against *Francisella tularensis*, the pathogen that causes the dangerous disease tularemia.

Delivery Makes the Difference

For the past eight years, Laboratory researchers have been developing and refining a platform for drug and subunit vaccine delivery based on tiny, disk-shaped structures called nanolipoproteins (NLPs). These closely resemble the good and bad cholesterol particles that move fats through the bloodstream and thus are unlikely to be flagged by the immune system as a foreign invader. Tests in mice and rats indicate that when antigens and other vaccine elements are attached to NLPs through a relatively simple process, the resulting vaccine elicits a stronger and more targeted immune response than when the components are administered separately. (See *S&TR*, September 2012, pp. 6–13.) In a collaborative effort, Livermore scientists have worked with numerous universities and companies to develop this approach.

"Subunit vaccines commonly consist of proteins from the pathogen of interest and molecules called adjuvants that further stimulate host immune response, together with materials to effectively deliver them in vivo," explains biologist Nicholas Fischer. "With the NLP platform, we can combine all these components in a single vehicle, so that all the pieces are delivered to every cell that internalizes the NLPs, all at the same time." He adds that NLPs, which range between 8 and 25 nanometers in diameter, are the perfect size to take advantage of natural paths into cells, particularly immune cells relevant to vaccine delivery. Moreover, the researchers have shown that NLP subunit vaccines can be freeze dried, stored for months at room temperature, and rehydrated without losing their potency.

So far the biggest test of the team's promising NLP platform has been the development of a subunit vaccine for *F. tularensis*, one of the most infectious bacterial pathogens in existence. (See *S&TR*, July/August 2010, pp. 15–17.) *F. tularensis* can infect an extremely broad range of hosts—more than 200 different animal species—and has a high mortality rate even at very low infectious doses—as few as 10 organisms. This rare but reemerging pathogen is most often spread by insect bites or inhaling contaminated particles but is also considered a potential bioterrorism agent. Despite decades of intensive research, no vaccine approved for widespread use yet exists. Leading live-vaccine candidates not only have significant side effects but also pose a substantial mutation risk because of their reliance on an attenuated bacterium with just a single genetic modification.

Pathogen Generates Complex Response

Vaccine development is rarely quick or easy, but the path to success for a *F. tularensis* vaccine has been strewn with a greater-than-usual number of obstacles. Research on such biothreat agents requires advanced containment facilities, of which a limited number exist. Furthermore, past studies have lacked standardization in such areas as the species and genetic strain of the test animal and the bacterial strain used, making it difficult for researchers to compare with and build on previous findings. For example, mice have been used extensively to study immune response to *F. tularensis* but are not an optimal choice because of their high susceptibility to all strains, even those that do not

cause illness in humans. Researchers now use a type of rat for such studies that better mimics human vulnerability to the pathogen.

According to Amy Rasley, Livermore biomedical scientist and *F. tularensis* vaccine development project lead, the biggest challenge is that biologists only partially understand the specific immune responses that a successful vaccine should provoke to this complex organism. She explains, "Most routine childhood vaccines focus on generating antibody responses. Past research suggests that antibody production is essential but not sufficient for protection against tularemia. Success rates for any vaccine drop when multiple immune responses are needed to combat the pathogen." Given the relative rarity of the illness in humans and the incomplete understanding of what constitutes a protective immune response, vaccine researchers must rely on rigorous experimentation and iteration, as well as patience.

Settling on the Right Combination

The initial Livermore approach was relatively simple. Building on previous *F. tularensis* studies, researchers created an NLP-based vaccine formula and used it to inoculate rodents. A few weeks later, the animals were exposed to a highly virulent, aerosolized form of *F. tularensis* to assess the degree of protection—whether the animals became ill, whether ailing animals recovered or died, and what the death rate was. Over a four-year period, the team tested dozens of antigen–adjuvant combinations, concentrations, delivery routes, and dosing schedules but never achieved more than 20 percent protection. "It was heartbreaking," confesses Rasley. "Something is definitely unique about *F. tularensis* that makes vaccine development exceedingly difficult."

However, a breakthrough came after Rasley and colleagues partnered with researchers at the University of New Mexico who were studying some of the same vaccine ingredients as the Livermore team for a different vaccine platform, one based on attenuated *Listeria monocytogenes* bacteria. Combining data from both teams led to the development of a multiple-antigen vaccine capable of stimulating strong antibody and T cell responses. The two teams found that using a combination of antigen types was critical—single antigens alone afforded only partial protection. Indeed, when Livermore researchers incorporated multiple antigens along with a promising adjuvant into their NLP platform, 100 percent protection against tularemia was achieved.

Interestingly, the team found that the survival rate was influenced by the vaccination route. All rats vaccinated through the nose survived, whereas only 90 percent of those vaccinated through muscle tissue did. Rasley speculates that generating a robust immune response in the lungs, as nasal delivery does, must be particularly important in combating tularemia. Rats vaccinated nasally also had less severe disease symptoms than their intramuscularly vaccinated cohorts.

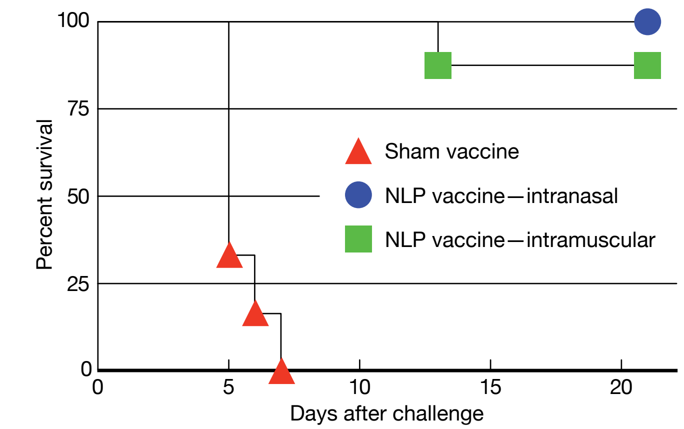

Test results indicate the efficacy of a tularemia subunit vaccine formula based on the NLP vaccine delivery platform. Up to 100 percent of vaccinated rats survived a bacterial challenge, depending on delivery method, as compared with the 0 percent survival rate of unvaccinated rats.

From What to Why

After identifying an effective vaccine, the team set out to understand why the formula works—a challenging investigation made somewhat easier by the subunit vaccine's relative simplicity and strict standardization. Fischer explains, "Determining what protects the host in a vaccine based on a whole, attenuated bacterium can be difficult because of the hundreds of synergistic components involved. In contrast, our formulation has only three components, so we can sort out the protection mechanism systematically. Once we understand the mechanism, we can see why it's working and identify correlates."

Correlates are additional standards by which the researchers can determine whether the vaccinated animals are protected, beyond simply testing whether the animal survives exposure. Identifying these correlates will be crucial for subsequent evaluations of how effective these vaccines are in other animal species, as well as in humans. Determining the protection mechanisms and correlates will occupy the next three to five years, according to Rasley. "We have a lot of important science to do," she states. The team is also reaching out to companies that specialize in producing vaccines at scale in accordance with reliable manufacturing processes. Although the Laboratory will not be commercially releasing the vaccine, demonstrating viable manufacturing at scale is another key step on the long path to vaccine approval and distribution—and the ultimate goal of robust protection against a dangerous pathogen.

—Rose Hansen

Key Words: antibody, antigen, attenuated bacterium, bioterrorism, *Francisella tularensis*, immune response, nanolipoprotein (NLP), subunit vaccine, tularemia.

For further information contact Amy Rasley (925) 423-1284 (rasley2@llnl.gov).

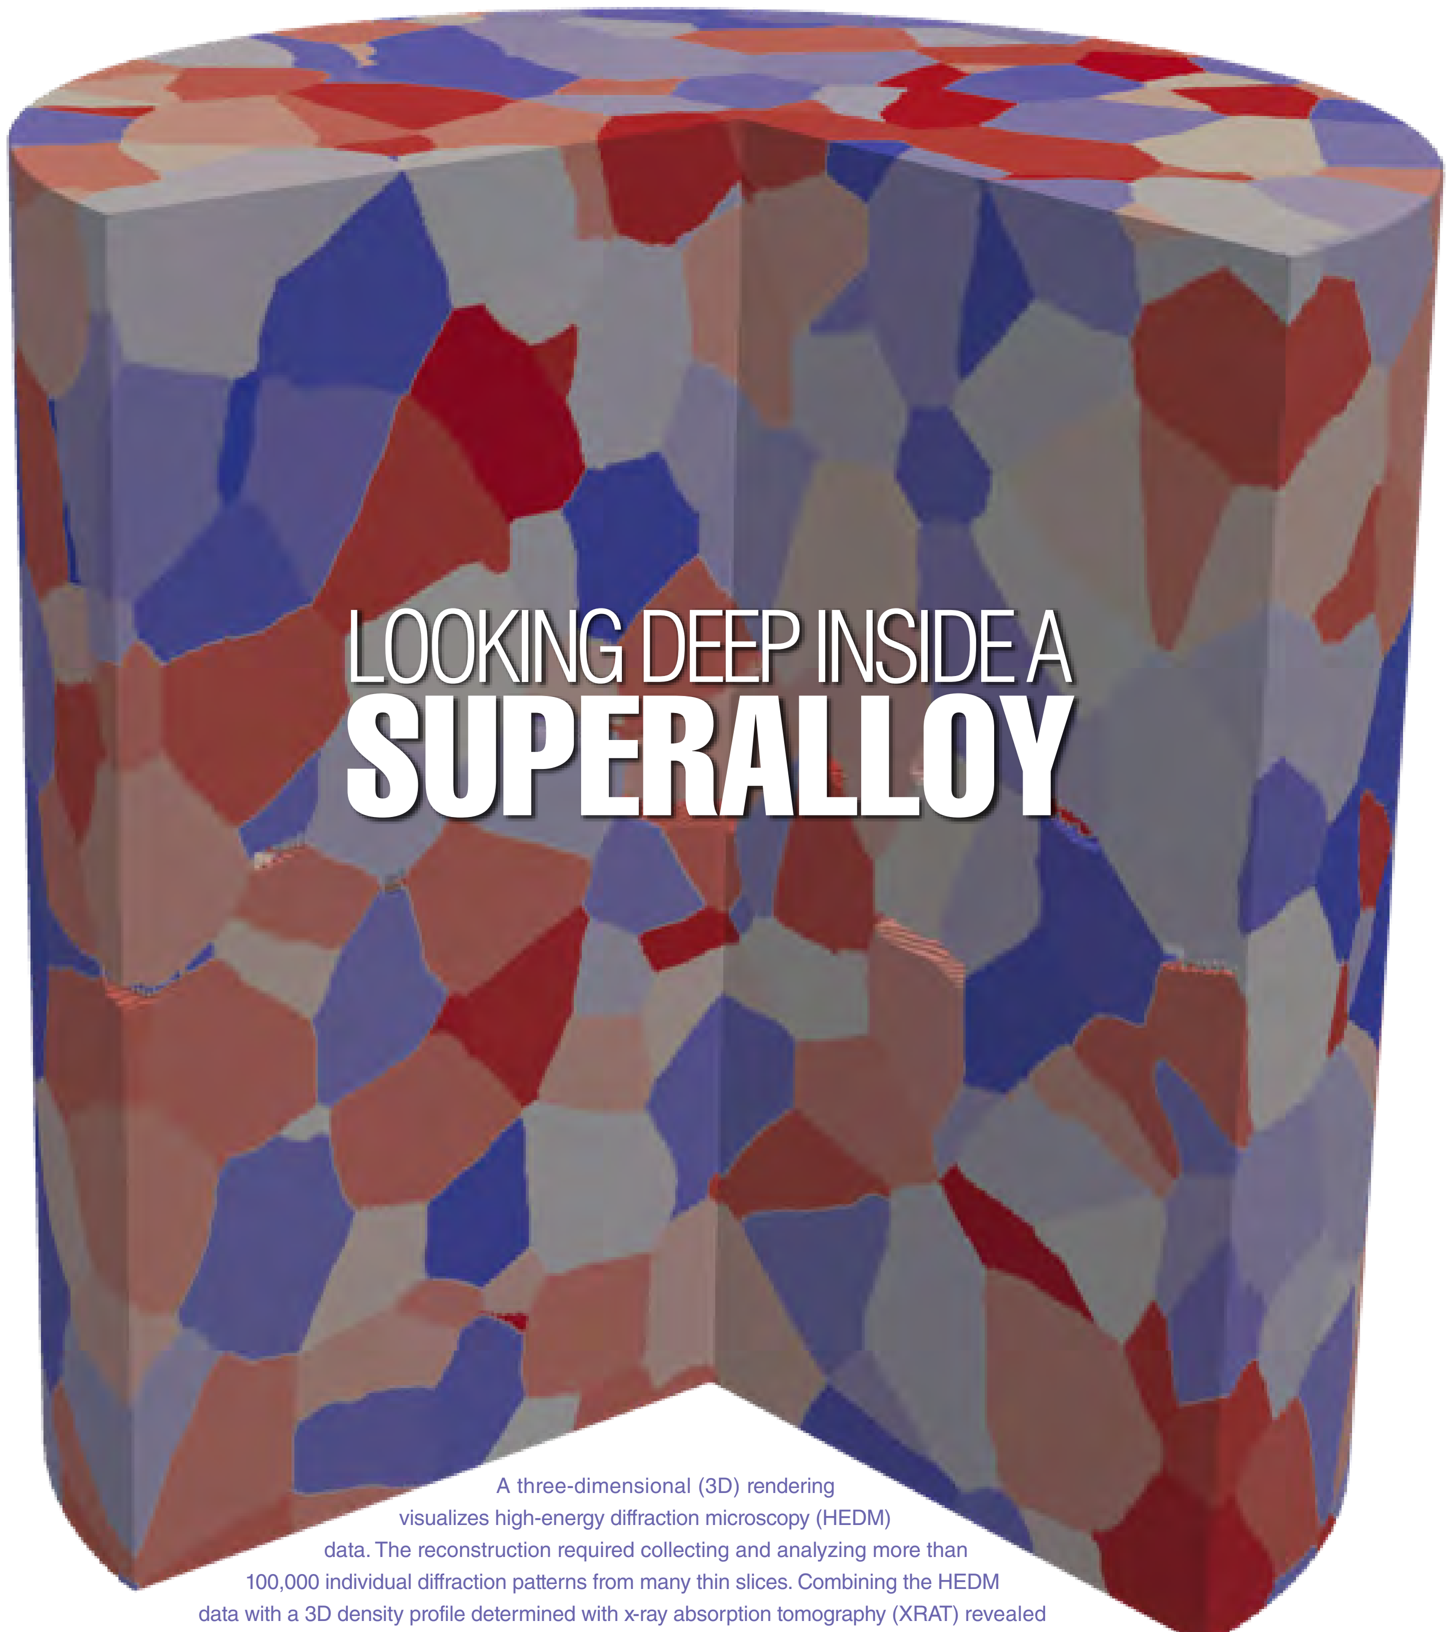

LOOKING DEEP INSIDE A SUPERALLOY

A three-dimensional (3D) rendering visualizes high-energy diffraction microscopy (HEDM) data. The reconstruction required collecting and analyzing more than 100,000 individual diffraction patterns from many thin slices. Combining the HEDM data with a 3D density profile determined with x-ray absorption tomography (XRAT) revealed the path of a crack and its relationship to grain boundaries (GBs) and other microstructure features.

DESPITE their monolithic appearance, metal parts used in various industries can consist of hundreds of millions of distinct microscopic crystalline particles called grains. Each grain is made of atoms arranged in a particular orientation, but the grains themselves are often randomly oriented relative to one another. Metals derive their strength and other mechanical properties from their internal structure, and the narrow spaces where individual grains meet—grain boundaries (GBs)—are where stress concentrates and cracks tend to form and spread.

Despite enormous strides in manufacturing methods to make metals stronger and more resistant to corrosion and fracture, parts subjected to harsh environments can still crack and ultimately fail, sometimes without warning. A leading cause of this failure is hydrogen embrittlement (HE), which is triggered by the entry and diffusion of hydrogen atoms, typically when the material is used in an aqueous or other corrosive environment. Even after a century of research, the physical mechanisms of HE are too poorly understood to predict HE-induced failure with high confidence.

A team of researchers, including Lawrence Livermore physicist Jonathan Lind, has applied nondestructive, synchrotron-based radiation to capture three-dimensional (3D) images of HE-induced microscopic cracks in metal. The images led the scientists to discover the relationship between the crystallographic character of metal grains and GBs and their susceptibility—or resistance—to HE. The team's research method, called near-field high-energy diffraction microscopy (HEDM), paves the way to improved predictions of HE based on the crystalline orientation of individual grains. The work, funded by the Department of Energy and the National Science Foundation, could also advance materials processing methods aimed at limiting cracks and thereby strengthening metal parts and extending their lifespans. (See *S&TR*, December 2014, pp. 16–19.)

Hydrogen Molecules Build Up Pressure

Lind explains that hydrogen entering a metal diffuses into the GBs, reducing ductility, toughness, and strength and often forming cracks that spread. These cracks can cause costly component failures requiring expensive repairs in a wide range of industries, including petrochemicals, nuclear energy, transportation, construction, and medical devices. Even strong, advanced alloys are susceptible to damage from HE. For example, nickel-base alloy 725, a so-called superalloy and the focus of the team's experiments, is designed for high strength and corrosion resistance in the oil-drilling industry, but nevertheless the material often exhibits cracks caused by HE. Lind says that improved strategies for predicting and preventing HE require deeper understanding of the 3D microstructural features that are most susceptible to this type of damage. Understanding the behavior of these features in the presence of hydrogen is key to improving lifetime predictions and designing HE-resistant microstructures.

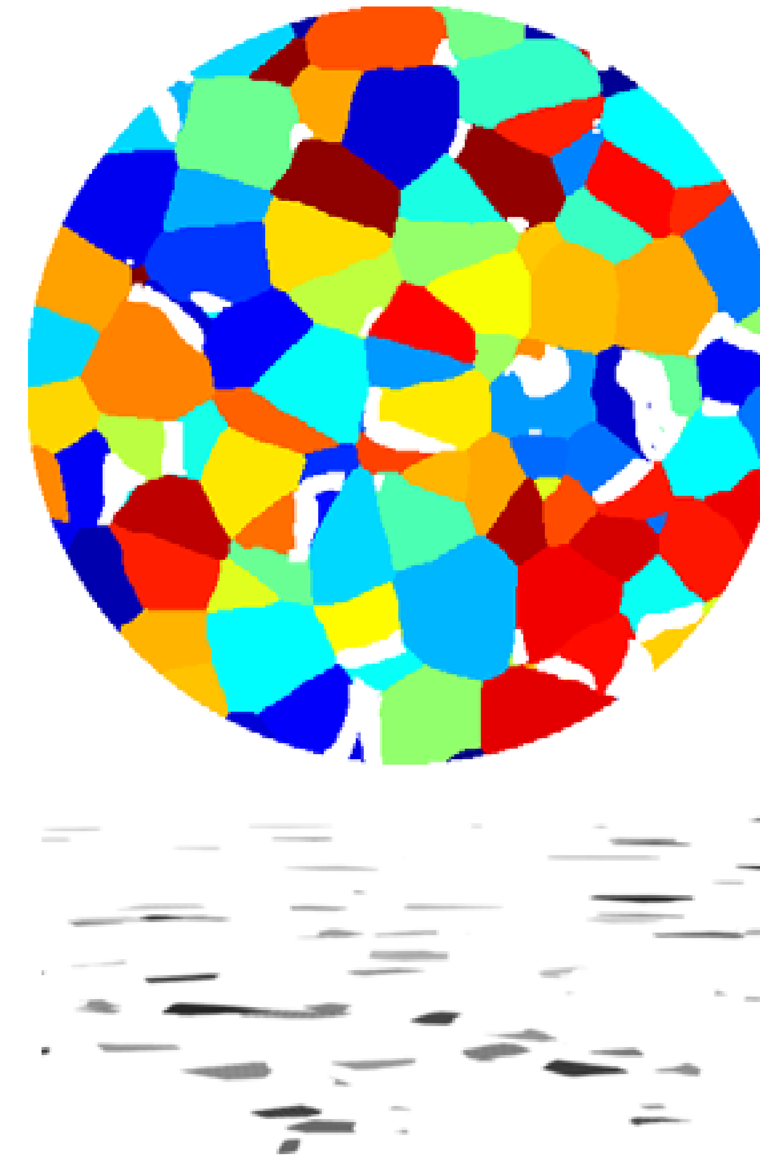

(top) A two-dimensional (2D) rendering is based on a reconstructed 1.2-millimeter-diameter slice of HEDM data. Colors represent single grains separated by GBs. Areas in white represent cracks. (bottom) This single x-ray diffraction pattern from an HEDM measurement was the basis for the 2D rendering. Individual diffraction spots capture individual grain shapes.

This effort brought together Lind with colleagues from the Massachusetts Institute of Technology (MIT), Argonne National Laboratory, Johns Hopkins University, Carnegie Mellon, and Texas A&M University. At Argonne National Laboratory's Advanced

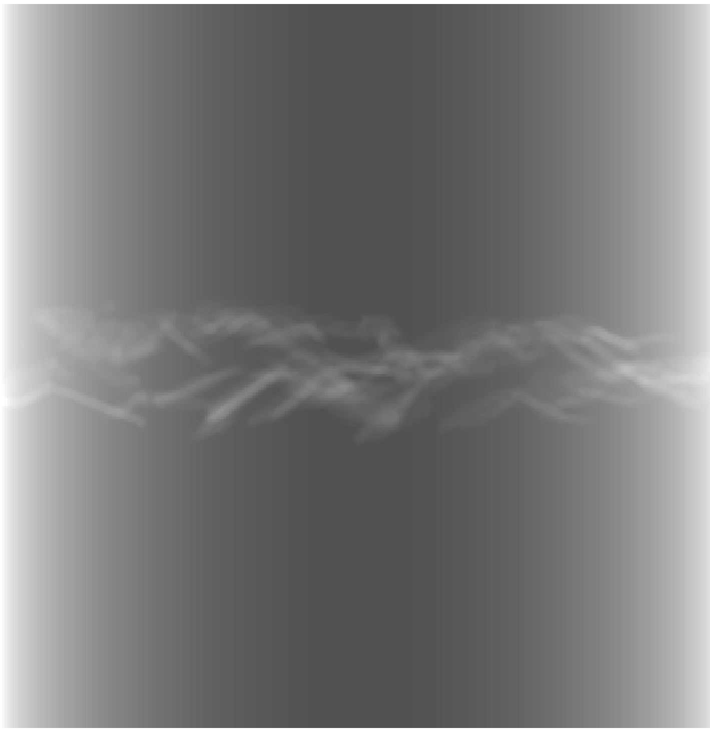

(above) In a single radiograph from XRAT analysis, the density contrast from a crack is readily apparent. (below) More than 1,000 radiographs were used to find the 3D shape of this crack in the reconstructed XRAT data.

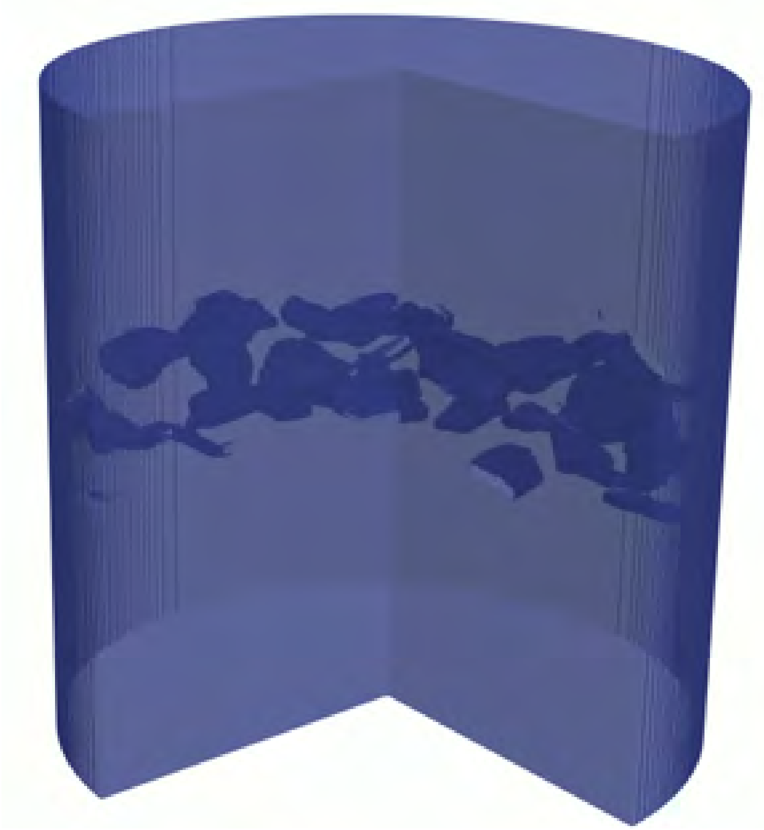

Photon Source (APS), HEDM was used to nondestructively image grain shapes and GBs of nickel-base alloy 725 with high-energy (50- to 100-kiloelectronvolt) synchrotron x-ray radiation from the 1-ID beamline. One of the few in the world specifically configured for HEDM measurements, the 1-ID beamline generates x rays by accelerating electrons to nearly the speed of light as they travel around a circular track. The resulting x-ray beam measures approximately 1 millimeter in width and penetrates a sample to a depth of several millimeters to several centimeters.

The researchers examined nickel-base alloy 725 specimens measuring 25 millimeters long and 3.8 millimeters in diameter. The material was electrochemically charged with hydrogen for seven days in solution, then slowly loaded in tension until failure, that is, the appearance of numerous cracks. One-millimeter-diameter cylindrical core samples were machined from each specimen's interior in accordance with the x-ray beam's size requirements. One sample contained the tip of a crack that extended across several grains with diameters of 90 to 125 micrometers. This crack tip—which became the main focus of the investigation—showed significant variation in orientation.

Two-dimensional (2D) HEDM diffraction patterns were obtained by passing an x-ray beam through the sample and capturing the diffracted x rays with a detector—a charge-coupled device camera paired with a scintillator that converts x rays to visible light. Because the x-ray beam is fixed in space, the sample was slowly rotated about an axis perpendicular to the beam over several minutes to bring all grains into multiple diffraction conditions. The procedure generated a sequence of 150 2D maps representing crystallographic orientations, which were then integrated into a 3D reconstruction of the microstructure.

Determining Crack Morphology

The team also took advantage of the beamline's capability for x-ray absorption tomography (XRAT). A form of computed tomography, XRAT allows a sample to also be scanned by HEDM without repositioning, thus enabling accurate alignment of the two data sets. With the XRAT data in hand, the scientists could link individual GB information collected by HEDM with the XRAT-derived 3D density information on crack morphology. Finally, Lind computationally reconstructed all the accumulated data from HEDM and XRAT into detailed, virtual 3D maps, or cross-sectional "slices" of the crack's twisting tip. This reconstruction was accomplished using software called IceNine, which was developed in part by former Livermore physicist Frankie Li. In addition, the final 3D reconstruction using IceNine yielded a virtual cylindrical form containing 50,000 grains and measuring 1.2 millimeters in height. The extremely detailed 3D map revealed the shapes, sizes, and crystallographic nature of grains along fracture surfaces.

"Researchers have been relying on 2D images for a long time," says Lind. "However, crack propagation is inherently a 3D problem, and so 3D information about cracks and how they travel through a metal is vital." By analyzing the 3D reconstruction, the team identified a set of so-called crack deflection events, where the crack deviates markedly from its original path despite having a less complicated path along which it could have propagated. In one example, a crack initially traveled along a single GB until reaching a "triple line," the intersection of three grains. The GB between two of the grains was well aligned with the crack plane and therefore appeared to be a likely path along which the crack could continue. Nevertheless, the crack proceeded along a different GB inclined at a high angle relative to the crack plane, which presumably required more energy.

Choosing the Road Not Taken

These surprising findings led the investigators to identify a class of GBs that strongly resist the propagation of HE-induced cracks—a GB where at least one grain exhibits a low Miller index, or a "boundary with low-index plane" (BLIP). The Miller index describes a plane of atoms in the lattice of a crystalline grain, with a low index indicating that the atoms in the plane are more tightly packed than those with a high index. Where grains meet, a crack propagating through a BLIP would therefore be energetically unfavorable. In effect, BLIPs deflect propagating cracks, thereby toughening the material and improving its HE resistance. Paraphrasing poet Robert Frost, Lind observes, "BLIPs represent the road not taken."

In the sample, the team identified 10 GBs that appeared to deflect cracks and prevent HE damage. Of these, nine fit the definition of a BLIP. By forcing the crack path to become more tortuous, BLIPs increase the total surface area created by the advancing crack, increasing the total work required and impeding its progression through a part.

The team's work is applicable to Lawrence Livermore's mission of stewarding the nation's nuclear stockpile. A key element of that mission is predicting the strength of materials under extreme conditions, and findings about BLIPs could advance simulations that use high-performance computing to better understand how metals behave under extreme pressure. (See *S&TR*, September 2018, pp. 12–15.)

The discovery of BLIPs and the development of 3D microstructure mapping could pave the way to improved predictions of the behavior of metals affected by HE. In addition, special material processing could create stronger GBs to hinder crack propagation in metal components regularly exposed to water or acid. Lind says, "If a metal could have many more BLIPs, its microstructures could better deflect and inhibit cracks and thereby significantly increase the lifespan of parts. Therefore, we want to develop methods for processing metals to have a higher fraction of BLIPs." The day may not be far off when stockpile stewards have further improved simulations and industries such as oil drilling have stronger, longer lasting metal parts.

—*Arnie Heller*

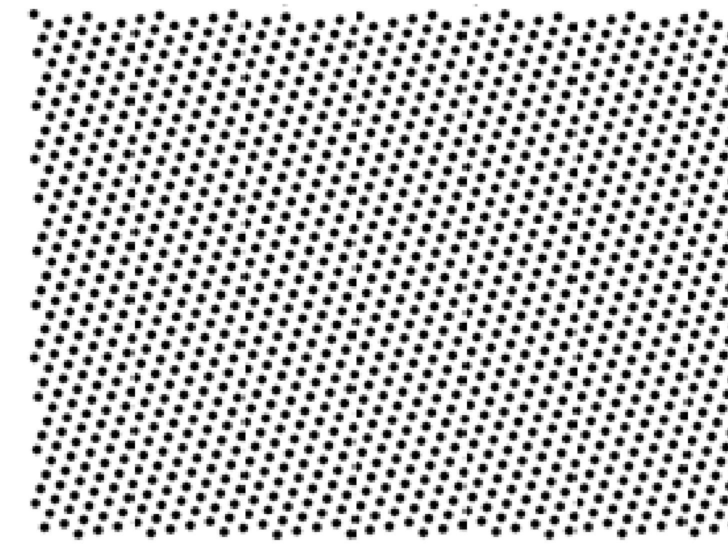

(above) A possible configuration of GBs with low-Miller-index planes—known as boundaries with low-index planes (BLIPs)—is shown. With their tight packing of atoms, BLIPs impede crack propagation, thereby strengthening the material. (below) In contrast, the loose packing of atoms seen in GBs with high-Miller-index planes present less of an energetic obstacle for crack propagation.

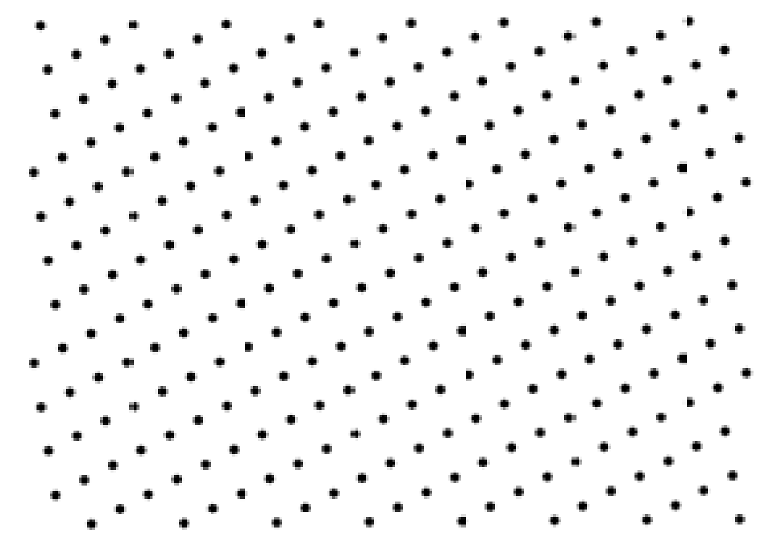

Key Words: Advanced Photon Source (APS), boundary with low-index plane (BLIP), grain boundary (GB), high-energy diffraction microscopy (HEDM), hydrogen embrittlement (HE), IceNine, Miller index, nickel-base alloy 725, x-ray absorption tomography (XRAT).

For further information contact Jonathan Lind (925) 423-8727 (lind9@llnl.gov).

Patents and Awards

In this section, we list recent patents issued to and awards received by Laboratory employees. Our goal is to showcase the distinguished scientific and technical achievements of our employees as well as to indicate the scale and scope of the work done at the Laboratory. For the full text of a patent, enter the seven- or eight-digit number in the search box at the U.S. Patent and Trademark Office's website (http://www.uspto.gov).

Patents

Cryogenic Pressurized Storage with Hump-Reinforced Vacuum Jacket
Salvador M. Aceves, Francisco Espinosa-Loza, Guillaume Petitpas, Vernon A. Switzer, Elias Rigoberto Ledesma-Orozco, Victor Alfonso Alcantar-Camarena
U.S. Patent 10,082,246 B2
September 25, 2018

Methods for the Selective Detection of Alkyne-Presenting Molecules and Related Compositions and Systems
Carlos A. Valdez, Audrey M. Williams
U.S. Patent 10,082,514 B2
September 25, 2018

Porous Materials via Freeze-Casting of Metal Salt Solutions
Michael Bagge-Hansen, Patrick G. Campbell, Jeffrey D. Colvin, Sergei Kucheyev, Thomas E. Felter
U.S. Patent 10,086,431 B2
October 2, 2018

Methods for the Selective Sequestration of Alkyne-Presenting Molecules and Related Compositions and Methods
Carlos A. Valdez, Alexander K. Vu
U.S. Patent 10,088,491 B2
October 2, 2018

Systems and Methods for Enhancing Optical Information
Peter Thomas Setsuda DeVore, Jason T. Chou
U.S. Patent 10,088,735 B2
October 2, 2018

Awards

The **German Research Foundation** (DFG) recently awarded Lawrence Livermore engineer **Brian Giera** with the **Mercator Fellowship**, a three-year visiting professorship at the University of Duisburg-Essen, one of Germany's largest universities. Giera's fellowship is a direct result of a keynote speech he gave at the 2017 International Electrophoretic Deposition (EPD) Conference in Gyeongju, South Korea. In the speech, he described a nanoparticle-based simulation of EPD, a process commonly used to coat objects with materials using electrical fields. Once his fellowship begins in spring 2019, Giera will collaborate with University of Duisburg-Essen Professor Stephan Barcikowski on a project to develop a method for uniformly coating neural implants with biocompatible metallic nanoparticles, adding his computational expertise to the effort.

As with the U.S. National Science Foundation, the DFG is a central, independent research funding organization. The latter foundation is supported primarily by the German federal and state governments. With the Mercator Fellowship's stipend, Giera will be able to travel to Germany several times each year, for a month or two at a time, to conduct research and give lectures.

Lawrence Livermore chemist **Dawn Shaughnessy**, whose team helped discover six new elements on the periodic table, has been elected a **fellow** of the **American Chemical Society** (ACS). At the Laboratory, Shaughnessy is the group leader for Experimental Nuclear and Radiochemistry and the principal investigator for the Heavy Element Group. She also led a group in naming heavy element 116, dubbed livermorium to honor the Laboratory and the city of Livermore, California. She and her team are currently working on a method to automate sample preparation and detection methods so that radiochemistry measurements will take one minute as opposed to up to five minutes when done by hand. The method, which will enable her team to analyze a single atom at a time, will also be applicable to other applications, including isotope analysis and nuclear forensics. The ACS fellows program began in 2009 to recognize and honor members for outstanding achievements in and contributions to science, the profession, and ACS itself.

The **Department of Energy's** (DOE's) **Exascale Computing Project** (ECP) has named **Lori Diachin** as its new **deputy director**, effective August 7, 2018. Diachin replaces Stephen Lee, who retired from Los Alamos National Laboratory. Since 2017, Diachin has been serving as the Computation Directorate's deputy associate director for science and technology at Livermore, where she has worked for 15 years. She previously worked at Sandia and Argonne national laboratories. She has held leadership roles in high-performance computing for more than 15 years. Her experience ranges from serving as director of the Laboratory's Center for Applied Scientific Computing to leading multilaboratory teams such as the FASTMath SciDAC Institute and directing DOE's HPC4Manufacturing and HPC4Materials programs.

ECP was launched in 2016 as a collaboration between DOE's Office of Science and the National Nuclear Security Administration to accelerate the delivery of an exascale computing ecosystem critical to DOE missions in national security, scientific discovery, and economic competitiveness. To achieve its goal of delivering exascale computing capabilities starting in 2021, ECP is developing strategies, aligning resources, and conducting research and design to deliver an ecosystem that includes mission-critical applications, a software stack, hardware architecture, and advanced system engineering and hardware components. ECP's collaboration includes experts from six core national laboratories—Argonne, Lawrence Berkeley, Lawrence Livermore, Los Alamos, Oak Ridge, and Sandia—along with representatives from industry and academia.

Abstract

Machine Learning on a Mission

More than just a buzzword, machine learning (ML) has become part of everyday life—from social media and online shopping to search engines and self-driving cars. ML uses computers to learn from and make predictions about data. A smarter machine can speed up processing and analysis times and improve the accuracy of identification and prediction tasks. Lawrence Livermore's national security mission and data-rich environment present ample opportunities for exploiting ML tools, often requiring new development beyond standard applications. In a broad range of applications, most research teams eventually seek ML-driven solutions, and many mission-critical programs already rely on these technologies. The Laboratory's go-to ML experts take a bidirectional approach, both advancing underlying theory and solving real-world problems. The former involves seeking mathematical solutions to optimize data sampling while exploring ML model interpretability through latent spaces and topological data analysis. Real-world applications include the multiscale modeling of interactions between cell membranes and specific proteins that induce many forms of cancer, as well as predictive image analysis for a range of screening scenarios—for example, medical diagnosis, airport luggage, and commercial truck cargo.

Contact: Peer-Timo Bremer (925) 422-7365 (bremer5@llnl.gov).

U.S. Government Printing Office: 2019/670-075-20002

Coming Next Issue

A Small Chip Promises Big Payoffs

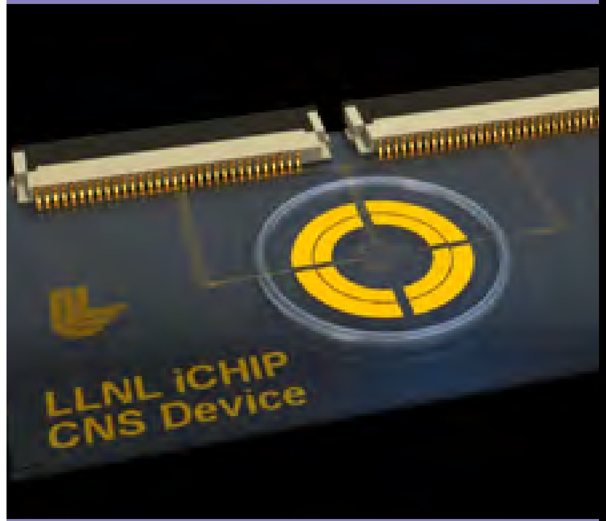

Livermore's brain-on-a-chip could advance understanding of how the brain functions, as well as aid the development of antidotes to toxic compounds, including chemical warfare agents.

Also in April/May

• A gamma-ray spectrometer being sent to space will provide insights into the formation of Earth and other rocky planetary bodies.

• Preparations for isotope harvesting at an advanced radioactive beam facility yield a surprising result.

• Using the National Ignition Facility, researchers create extreme x-ray and neutron conditions to test the nation's nuclear stockpile.

www.llnl.gov

Science & Technology Review
Lawrence Livermore National Laboratory
P.O. Box 808, L-664
Livermore, California 94551