

COGNITIVE SIMULATION SUPERCHARGES SCIENTIFIC RESEARCH

Adding machine learning and other artificial intelligence methods to the feedback cycle of experimentation and computer modeling can accelerate scientific discovery.

COMPUTER modeling has been essential to scientific research for more than half a century—since the advent of computers sufficiently powerful to handle modeling’s computational load. Models simulate natural phenomena to aid scientists in understanding their underlying principles. Yet, while the most complex models running on supercomputers may contain millions of lines of code and generate billions of data points, they never simulate reality perfectly.

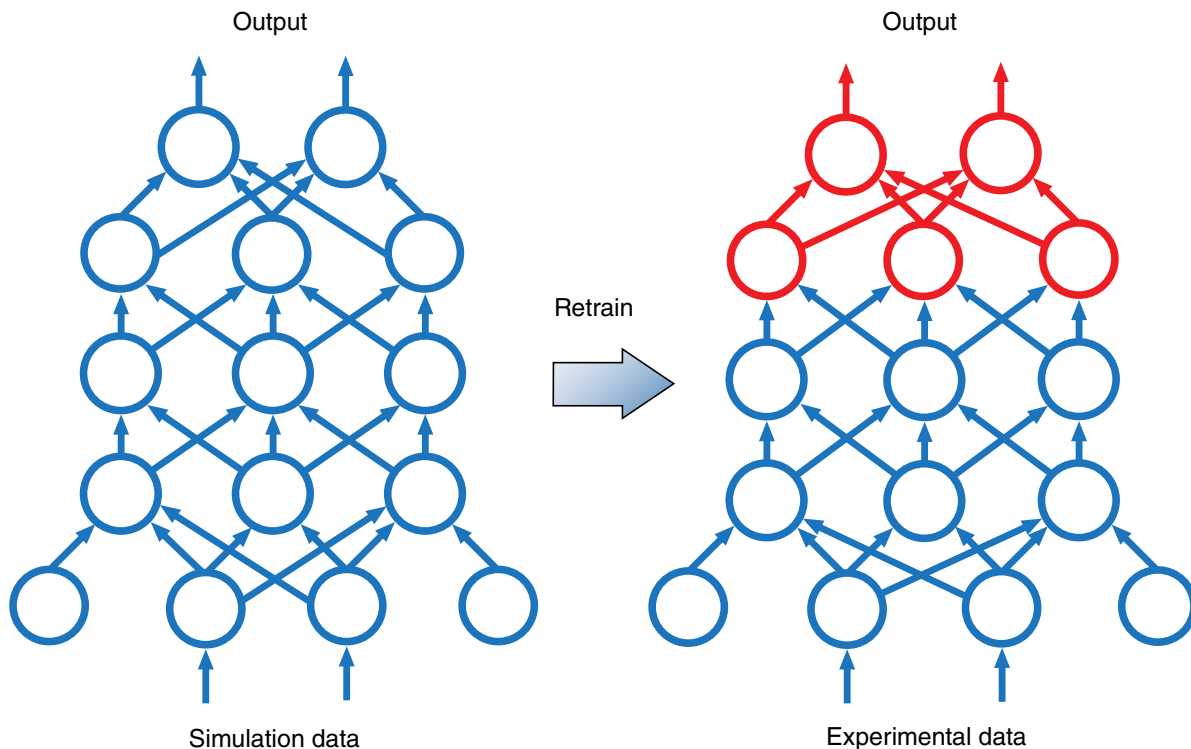
Experiments—in contrast—have been fundamental to the study of natural phenomena from science’s earliest days. However, some of today’s complex experiments generate too much data for the human mind to interpret, or they generate too much of some data types and not enough of others. The inertial confinement fusion (ICF) experiments conducted at Livermore’s National Ignition Facility (NIF), for example, generate large volumes of data; but fusion physics is complex, and connecting the underlying physics to the available data is a difficult scientific challenge. Researchers increasingly face the problem of how to process data so that the underlying physics emerges into light.

To improve the fidelity of complex computer models, and to wrangle the growing amount of data, Livermore researchers are developing an array of hardware, software codes, and artificial intelligence (AI) techniques such as machine learning (ML) they call cognitive simulation (CogSim). Researchers will use CogSim to find large-scale structures in big data sets, teach existing models to better mirror experimental results, and create a feedback loop between experiments and models that accelerates research advances. CogSim’s goal is ambitious—to transform itself into a fourth pillar of scientific research, joining the three pillars of theory, experiment, and computer modeling, as tools of discovery.

A New Feedback Loop

“Our wonderful simulations, as good as they are, are not perfect representations of our experiments,” says Brian Spears, a physicist and leader of the Laboratory

In a deep neural network embedded within a modeling software code, simulation data flows through linked nodes (shown as circles) that numerically represent underlying physical behavior and train the software on the simulation data (left). The software is then retrained with experimental data, producing a model that both represents the physics and replicates experimental data (right). Red layers represent the retrained data subset.



Director's Initiative in CogSim. "We make approximations in these models. They have shortcomings." Models cannot incorporate physics that is not yet known. Scientists use data from experiments and simulations, evaluate these results together and incorporate their learning into the model to help design the next experiment. The model's accuracy at representing the experimental results is a measure of how well the models replicate the science behind the experiment. "The process has worked well for 50 years," Spears notes. Now, however, AI-ML tools offer a way to better analyze the results, incorporate them more completely into the modeling framework, and use the "retrained" model to improve the design of experiments with the goal of accelerating discovery.

A deep neural network (DNN) is a sequence of computing layers that transforms the inputs until they produce the desired output. A node, also called a

neuron, is the basic unit of computation in a neural network. It accepts external inputs from other nodes and applies a mathematical function that results in an output. "A distinctive layer in the interior (of the DNN) we call the latent space," explains Spears. "This space is the most distilled representation of the physics in the data that we can get our hands on. In the latent space, we've preserved all the correlations in the physics that we can use to reduce the amount of data we need to accelerate the training." With the input of the large data set through the DNN, the modeling software is now trained on the simulation data. Then, researchers retrain a portion of the DNN on experimental data.

"We can train AI's deep-learning models on our simulations to make a perfect surrogate that knows exactly what the model's software code knows. However, the model is still not good enough because of the shortcoming of the code to predict the experiment," explains

Spears. "We then retrain part of this model on the experimental data itself. Now we get a model that knows the best of both worlds. It understands the theory from the simulation side, and it makes accurate corrections to what we actually see in the experiments." The new improved model is now a better reflection of both the theory and the experimental results.

Applying CogSim to a wide range of research could benefit many fields. "At its highest level, we can build cognitive simulation into the Laboratory's research strategy, and all of its missions," says Spears. "We're already applying it to inertial confinement fusion and the Laboratory's stockpile stewardship mission, as well as in bioscience research. But we're driving it into other areas."

Laboratory researchers are developing three technologies to transform how scientists predict the results of experiments through CogSim. The first is refined models that fully use available data

sets, estimate uncertainty, and are further improved through exposure to experimental data.

Software tools to develop and guide application of those models are the second need. These tools include workflows that guide simulations, learning frameworks that scale to the largest platforms at Livermore, and an environment in software that allows users to work productively without the hurdle of building the software stack and understanding the interface.

The third need is for computational platforms to support these tools. The Laboratory anticipates that its first exascale computer, El Capitan, arriving in 2023, will be a driving force of CogSim.

Improving ICF Experiments

Laboratory researchers already use Livermore-developed CogSim tools to design new ICF experiments to achieve nuclear fusion ignition. During a NIF shot, as many as 192 lasers fire into a hohlraum, a hollow cylinder open at the ends that holds a 2-millimeter-diameter fuel capsule

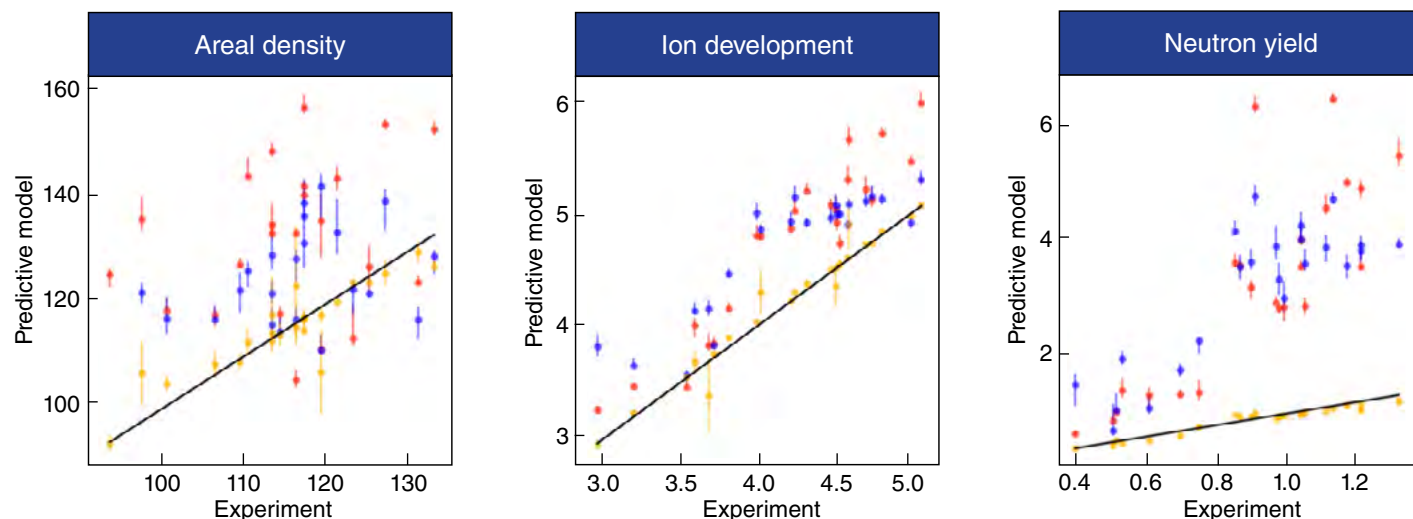
containing a mixture of the hydrogen isotopes of deuterium and tritium.

The lasers strike the interior walls of the hohlraum, converting their ultraviolet-wavelength light into x-ray wavelengths, bathing the fuel capsule in an intense burst of energy. The capsule shell explodes, compressing the fuel inside into a pinpoint-size area of high-energy-density plasma. The implosion causes hydrogen atoms to fuse, releasing energy. According to the widely used National Academy of Sciences definition, fusion ignition takes place when the energy output from the shot is greater than the energy input. NIF researchers have developed software to model these experiments. The model inputs include factors such as the thickness of the fuel capsules, the geometry of the hohlraum and size of its laser entry holes, the fuel mix and gas fill inside the capsule, and the energy and pulse shape of laser shots. Physical parameters affect the energy output, and researchers measure quantities such as the intensity of emitted radiation, number of neutrons, and the evolution of ion temperature inside the

plasma with time to see how their design affected the shot results. By modeling experiments with different inputs, the researchers can try to increase critical outputs that approach fusion ignition and avoid spending time and money on setups that won't improve the outcome.

A CogSim method called transfer learning is helping researchers improve ICF models. "Transfer learning is a two-step process," explains physicist Kelli Humbird. "First, we need to train a neural net from scratch on a really large data set relevant to the problem. Neural nets are data-hungry, so they need many examples of how to solve a given problem to learn effectively. Second, we take a small data set that we care about, that should be related to the large set, and partially retrain the neural net on that data. The large data set should give the net a general idea of the problem it's trying to solve. The small data set fine-tunes its predictions on the more targeted task."

ICF simulations can be used to create large data sets. "This gives the model a good idea of how ICF implosions change



Blue dots in the graphs represent low-fidelity model predictions of an inertial confinement fusion experiment model before the experiments took place; red dots represent a higher fidelity model's pre-shot predictions. Yellow dots represent the cognitive simulation (CogSim) model's ability to predict the results of each experiment based on transfer learning from the aggregate set of collected experimental data.

as you change input conditions—the target, the laser pulse, and so on,” says Humbird. “However, our simulations are not perfect—physics approximations mean that what our simulation predicts is sometimes not what we see in the experiment. So, we can take our neural net with ‘general knowledge’ about ICF, and our small set of experimental data, and retrain just a portion of the model with the experiments. Now, our neural net has basically learned how to take the simulation predictions and adjust them to


be more in line with what we really see in experiments.”

In an experiment applying CogSim to ICF experiment simulations, Humbird modeled a series of laser shots conducted at the University of Rochester’s Laboratory for Laser Energetics. Before the experiments, they modeled measurable parameters such as the areal density of a shot’s plasma cloud, the ion temperature, and the neutron yield and then used the data from a portion of the shots to retrain their model using

a DNN. The retrained model more correctly predicted the results of the other experiments. Their paper, “Transfer Learning to Model Inertial Confinement Fusion Experiments,” reported in *IEEE Transactions on Plasma Science* won the 2022 Transactions on Plasma Science Best Paper Award from the IEEE Nuclear and Plasma Sciences Society. The model in this proof-of-concept experiment used just three hidden layers (DNN layers between the input and output layers) with a few dozen neurons, nine inputs, and five outputs. In fall 2022, *HPCwire* magazine awarded its Editor’s Choice award to the Livermore team for Best Use of HPC in Energy, recognizing the application of CogSim to ICF research.

The group also ran simulation and experimental data from ICF shots at NIF using transfer learning, and a larger amount of data. The results indicate that the DNN-trained model provided the best match with the experimental results. “We’ve been making predictions with the NIF model routinely and updating it as we acquire more data. We’re observing that it generally improves over time, and we’re hoping to add new diagnostic data, such as capsule quality, and other metrics we think impact performance,” Humbird notes.

In August 2021, a NIF shot yielded an output of about 70 percent of the input energy—the threshold of ignition. (See *S&TR*, April/May 2022, pp. 4–15.) A second NIF shot in September 2022 that applied record-high laser energy produced about 1.2 megajoules (MJ) of fusion energy yield, compared with the 2021 experiment that produced 1.35 MJ. Later in the fall, CogSim techniques predicted that a specific experiment conducted at NIF had a greater than 50 percent probability of surpassing the ignition threshold. The model was correct; the record yield observed in the December 2022 experiment fell within the range predicted by the CogSim



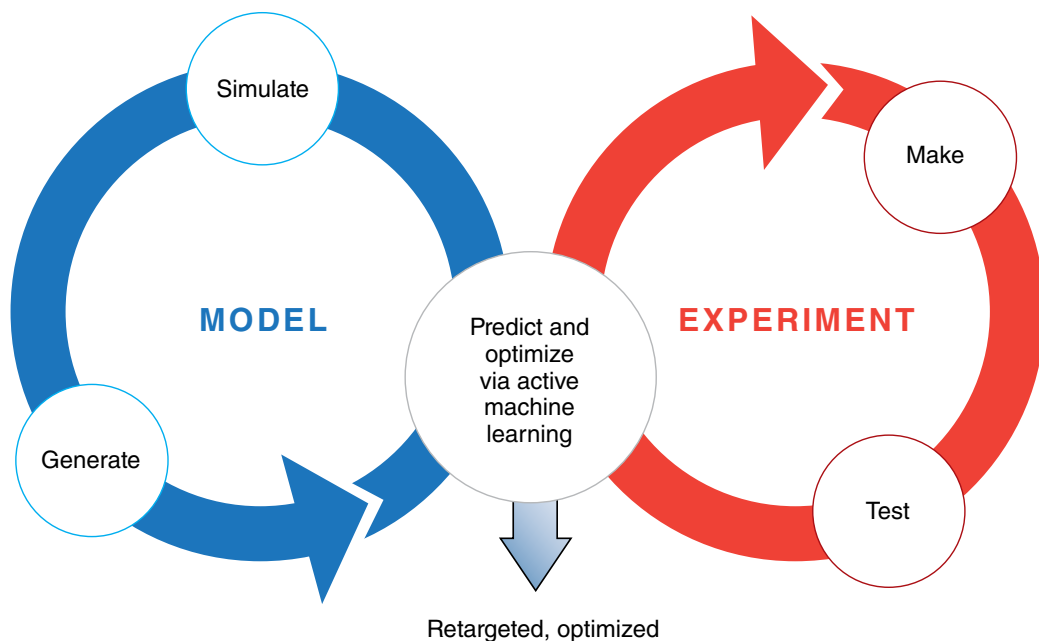
Livermore researchers trained a model to design mutations to a SARS-CoV-1 antibody that enabled it to bind to and neutralize SARS-CoV-2, the virus that causes COVID-19. Illustrated is a structural model of the receptor binding domain (blue) of SARS-CoV-2 spike protein in complex with modified m396 antibody (red) and designed antibody mutations (green).

model. “The ability to predict the outcome, with credible uncertainties, of such a significant experiment demonstrates great improvements in simulation capabilities and physics understanding,” says Humbird. “The growing use of CogSim techniques strengthens predictive capability, enabling a push into high performing regions of the design space at NIF.”

Faster Discovery for Bioscience

CogSim tools can work across a range of scientific endeavors. Laboratory bioscience researchers have been applying CogSim to produce predictive models for biology since the 2000s. One example is the Accelerating Therapeutic Opportunities in Medicine (ATOM) consortium, a partnership of GSK (formerly known as GlaxoSmithKline), Lawrence Livermore, Frederick National Laboratory for Cancer Research, and the University of California at San Francisco. ATOM develops a preclinical design and optimization platform to help shorten the drug discovery timeline. The consortium uses ML to predict the properties of proposed drug molecules, and screens them virtually for safety, pharmacokinetics, manufacturability, and efficacy.

Jim Brase, deputy associate director for Computing, leads Livermore’s ATOM efforts, and the Laboratory’s work applying HPC to life science and biosecurity. “There isn’t enough data for some parameters to build the high-fidelity model we want,” he says. “The chemical space is so wide open. Current models can’t predict a protein structure to have the properties that enable it to, for example, penetrate the blood–brain barrier. So, we have to build a machine-learning model to teach the simulation how a protein penetrates the barrier and use this understanding to predict the proteins that might do a good job at the task.”



In the biosciences, a cycle of simulating mutated antibody reactions with a virus and improving mutations through simulation coupled with synthesizing and testing antibodies can speed antibody developments from years to months or weeks.

Another Laboratory effort to apply CogSim to countering biothreats was accelerated by the COVID-19 pandemic. (See *S&TR*, June 2021, pp. 4–11.) Daniel Faissol led a team to design antibodies capable of binding to SARS-CoV-2 using ML algorithms and Livermore’s high-performance computers. The team took antibodies known to be effective against SARS-CoV-1 and retargeted them by proposing mutations of amino acid sequences on the antibody they thought would work against SARS-CoV-2. ML algorithms helped identify promising sequences, and high-fidelity simulations tested how well the antibodies bound to a SARS-CoV-2 antigen to inactivate it. Faissol, who leads Livermore’s program in Computational Design of Antibodies, says, “We did this with three different SARS-CoV-1 antibodies and successfully retargeted them to neutralize SARS-CoV-2, sometimes

maintaining full potency against SARS-CoV-1.”

When the COVID-19 pandemic hit, the team had been working to extend the protective effects of a meningitis-B vaccine by modeling the meningitis virus’ antigen proteins. To combat SARS-CoV-2, Faissol says, “We flipped this model around—we modeled different candidate antibodies to SARS-CoV-2, while keeping the antigen the same. It was a different approach compared to what others were doing. The method is very fast and can potentially result in broadly protective antibodies against COVID variants.” Laboratory development of antibodies normally takes years. “Because we’re doing this computationally, we can do this work rapidly, in principle. We didn’t have to wait for a patient who was infected with and recovered from both SARS-CoV-1 and -2 to design a protective antibody,” he notes.

Later, the team was asked by Department of Defense (DOD) sponsors to redesign a clinically used SARS-CoV-2 antibody—part of AstraZeneca’s antibody drug product for COVID-19 that suffered significant potency losses against Omicron variants. “We performed a lightning computational sprint over two weeks that successfully identified antibody variants with potency that bound Omicron-variant antigens while maintaining full potency against previous strains,” says Faissol. “One of the things that distinguishes our work is co-optimization for multiple antibody properties, such as stability, together with binding to multiple viral variants. The Omicron antibody binds with multiple Omicron antigens—we can optimize it for manufacturability at scale.”

Thanks to this work, DOD has funded a six-year effort to further develop this CogSim-based optimization platform called Generative Unconstrained Intelligent Drug Engineering (GUIDE). The goal of the GUIDE program is to

develop the predictive tools, data, and experimental capability to enable rapid response to new pathogens and to design broadly protective antibody and vaccine candidates against entire classes of viral pathogens. The platform will also help accelerate the time required to design robust medical countermeasures to days or weeks, and it will provide capabilities for early analysis of emerging biothreats. This work aligns with the Laboratory’s mission focus area in bioresilience—the use of data and CogSim tools to greatly improve the response to biothreats. The Laboratory’s cooperative work with AstraZeneca is just one example of a collaboration furthering CogSim.

Collaboration Speeds AI-ML


In December 2021, the Laboratory announced the founding of the Artificial Intelligence Innovation Incubator (AI3), a collaborative hub aimed at uniting experts in AI from Lawrence Livermore, industry, and academia to advance AI for large-scale scientific and commercial

applications. The Laboratory is driving AI3 to advance scientific solutions using AI tools—building partnerships with Google, NVIDIA, Hewlett Packard Enterprise, Advanced Micro Devices, and others. Tools developed can be brought back to the Laboratory to advance its missions.

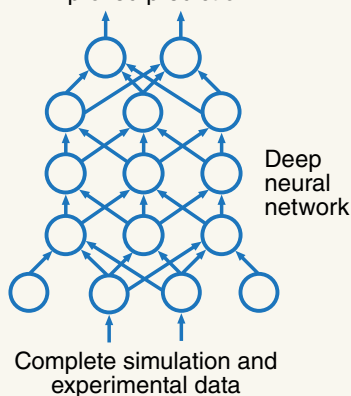
Laboratory scientists will apply CogSim tools to stockpile stewardship and other missions through discovery, design exploration, manufacturing and certification, and deployment and surveillance. Laboratory researchers will use CogSim tools to design new molecules and materials vital to national security priorities. Design exploration tools will help them accelerate development of technologies used in the nuclear stockpile. Manufacturing tools will help increase the speed and efficiency of manufacturing parts and reduce materials waste.

With these tools, Livermore researchers will answer the question, “How can we design a material that performs well in the component we’re building and is also easy


Traditional Pillar
High-performance computing



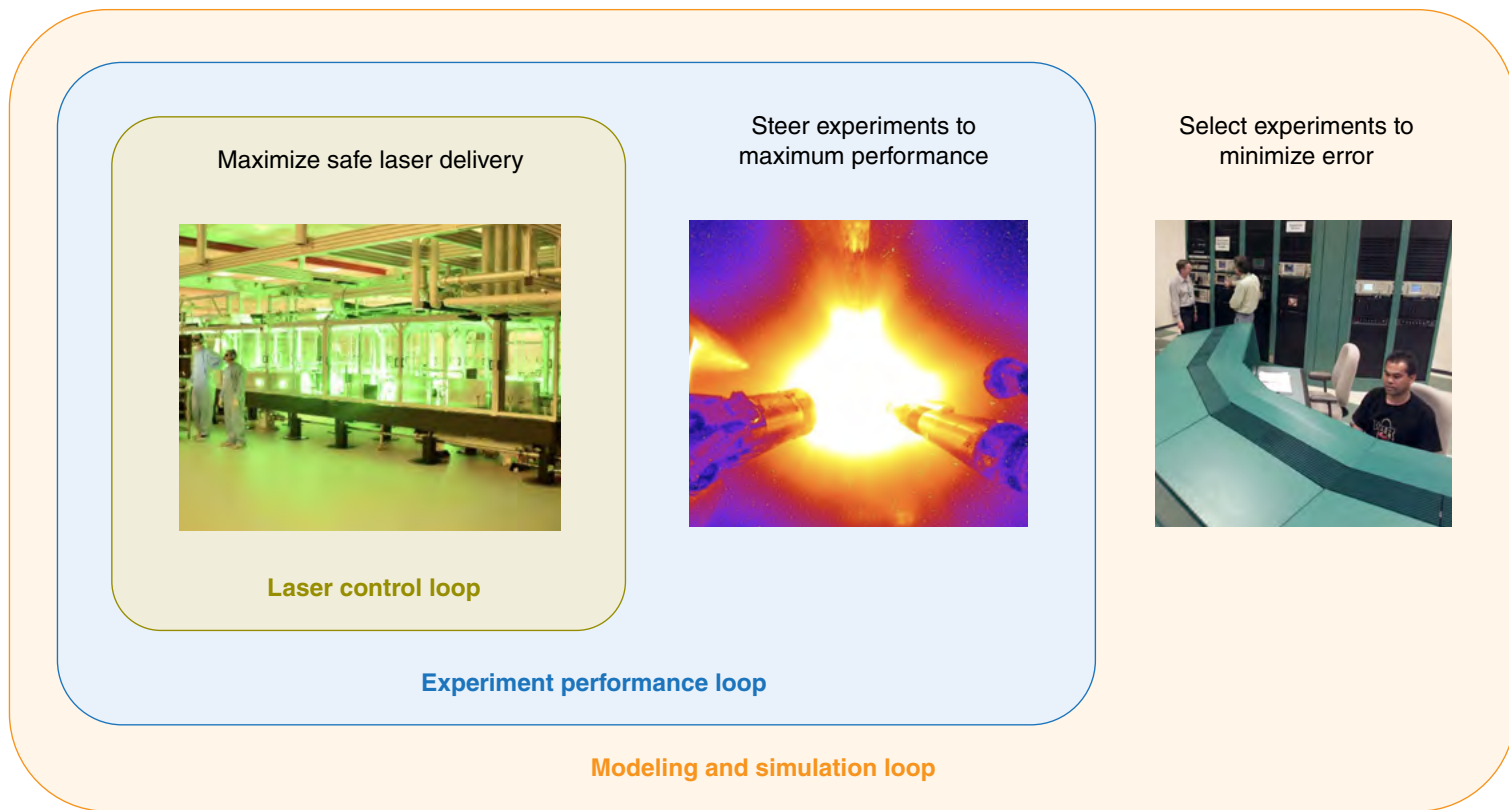
New Pillar
Cognitive simulation
Improved prediction



Traditional Pillar
Large-scale experiments



Adding a new pillar of cognitive simulation, researchers use machine learning to apply the results of a model, such as those from the hydrodynamic simulation code HYDRA (left), to results of an experiment, such as an x-ray image from the National Ignition Facility (NIF) (right). Running the improved model helps guide the design of the next experiment.



Livermore is working with chip manufacturer NVIDIA to develop CogSim-based software to control and steer laser experiments for maximum performance.

to manufacture?” Researchers will develop a new material, design a part for the stockpile or for ICF experiments and then apply CogSim tools to develop efficient manufacturing processes. AI offers the ability to pull these tasks closer to each other. Next, in considering performance over the component’s lifetime, researchers will answer questions such as, “How will systems change over their lifetime? What’s the impact of aging? Are the materials compatible with each other?” Says Spears, “Cognitive simulation tools will ensure that we can design for the lifetime of the technology so that, for example, a material won’t degrade early in a device.”

Self-guided Laser Experiments

Work is also underway to use CogSim to help drive the future of high-power laser

experiments for ICF and basic science. Laboratory researchers are working with chip manufacturer NVIDIA to develop a CogSim-based automated control system for laser experiments. NVIDIA’s research on self-driving cars has produced hardware and software that the Laboratory is adapting in partnership with the company to build a self-driving laser control system. “We’ll be able to control laser experiments sufficiently well to perform one every minute, and we know how to scale this rate up to one shot per second,” says Spears. CogSim embedded in a computer-controlled system could quickly reset the parameters of a laser shot to achieve a desired result. The goal is to accelerate discovery through high-energy, high-repetition rate laser-driven experiments (See *S&TR*, July 2021, pp. 20-23.)

In just a few years, CogSim has gone from small-scale studies to a promising tool that is poised to broadly impact Livermore’s missions and research. New discoveries will not trail far behind.

—Allan Chen

Key Words: Accelerating Therapeutic Opportunities in Medicine (ATOM) consortium, artificial intelligence (AI), Artificial Intelligence Innovation Incubator (AI3), AstraZeneca, cognitive simulation (CogSim), COVID-19, deep neural network (DNN), El Capitan supercomputer, exascale computing, hohlraum, inertial confinement fusion (ICF), machine learning (ML), National Ignition Facility (NIF), NVIDIA, Omicron variant, SARS-CoV-1, SARS-CoV-2.

For further information contact Brian Spears (925) 423-4825 (spears9@llnl.gov).