

MACHINE LEARNING POINTS TOWARD NEW LASER TARGET DESIGNS



WHEN the Trinity supercomputer at Los Alamos National Laboratory was first coming online, calls went out for research projects that would test—and potentially break—the new system. Researchers from Lawrence Livermore answered the call, and their work with Trinity and machine learning could disrupt 40 years of assumptions about inertial confinement fusion (ICF).

“The theory of ICF was all done with pencil and paper, assuming a spherical implosion,” says design physicist Luc Peterson. “In many studies, if your implosion isn’t spherical, you’re not getting as much energy out of it as you could.” ICF implosions at Livermore’s National Ignition Facility (NIF) are

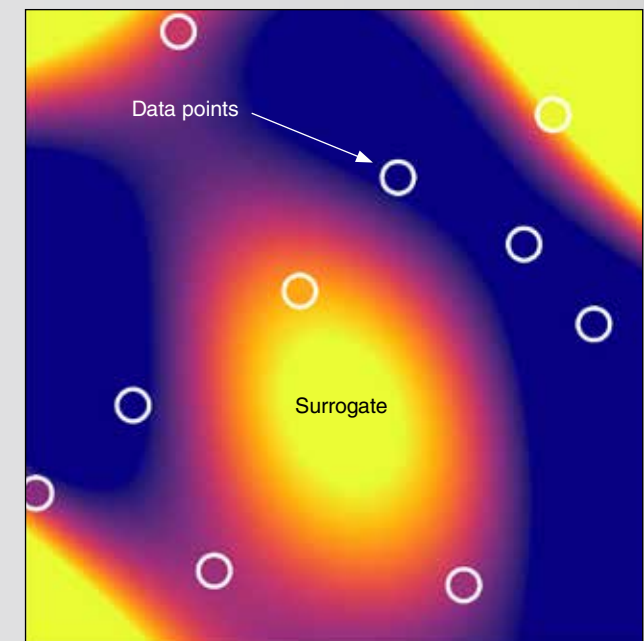
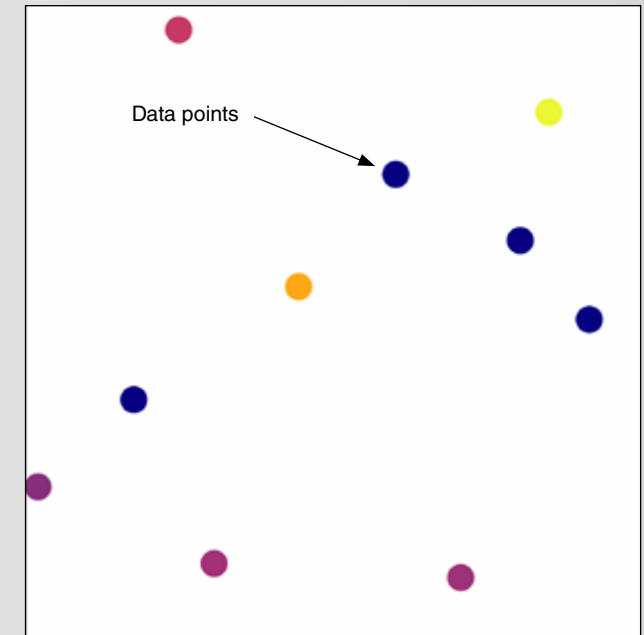
aimed at a spherical target housed in a cylindrical hohlraum, which creates an asymmetrical, preferred axis. To combat this tendency, Peterson was tasked with what he calls an impossible job: either make the implosions more round or create an implosion robust enough to withstand the inherent asymmetries and still achieve high energy yield. “I listed all the ways that NIF could possibly implode something asymmetrically,” Peterson says. “I got a very large number of parameters and realized that to check all the different combinations, I would need to run many simulations—more than had ever been done before.”

The nine parameters included various asymmetries, drive multipliers, and gas fill densities—all factors that affect the

quality of target implosion. Simulating all the permutations would produce 5 petabytes of raw data, which is close to the current limit for Livermore’s parallel file systems. Steve Langer, Peterson’s colleague and a fellow Laboratory design physicist, heard about the effort to map out all nine parameters and conceived of a way to help.

From Supercomputer to Server Farm


Langer’s idea involved Trinity, then a brand-new Cray XC40 system. Typically, before transitioning a new computer to classified work, a national laboratory holds an open-science period where researchers can “kick the tires” of the new system



(top) Simulations on the Trinity supercomputer produced approximately 60,000 data points, which were then used to train a machine-learning model. (bottom) The model predicted all the points between the simulations to produce a surrogate model across nine parameters, which can be represented by the gradient in any two dimensions of the nine-dimensional space.

by running unclassified experiments. Laboratory technicians can also consult with the computer's vendors as the experiments run and discover ways to fine-tune the system. Langer's plan was to process their raw physics data on the fly, analyzing and deleting files while they were being created, instead of saving all the data. Peterson and Langer pitched their big-data physics simulation proposal to Los Alamos, and a collaboration was born. "We knew we would have to do some distillation to even store the results on disk, which prompted us to create this on-the-fly, in-transit system," says Peterson. "We developed a system to perform the filtering while the simulations are running. The approach is like filling up a bucket with water while making a hole in the side to drain the bucket so it doesn't overflow."

Their project essentially turned Trinity—then a 8.1-petaflop (10^{15} floating-point operations per second) supercomputer designed to run one large simulation at a time—into a giant "server farm" capable of running several thousand simulations at once. The approach was not only necessary for the project but also worthy of the new computer's open-science challenge, stressing Trinity in new, often unforeseen ways that sometimes affected other users. One surprised Los Alamos employee sent out a midnight email asking whether someone was performing large data transfers that had lowered data rates to only 17 gigabytes per second for codes that normally achieved more than 600 gigabytes per second. "Los Alamos put me on



speed dial for what I did to their poor machine," says Peterson. The filtering system managed to trim the expected 5 petabytes of raw data down to 100 terabytes, but transferring the data between the two laboratories still took several months. "We joked that it would actually be faster to rent a van and drive across the desert with a bucket of USB drives," he adds.


Rise of the Machines

Generating the data was only half of the challenge. The next step was to analyze the data and search for robust designs. However, searching through all the simulations was not sufficient. "We have approximately 60,000 data points, which sounds substantial, but when you consider nine-dimensional space, it's actually pretty sparsely sampled," says Kelli Humbird, a Livermore Graduate Scholar who helped Peterson study the data. "We wanted an algorithm that would interpolate between the points and connect the dots so we could approximate the results of simulations anywhere in the nine-dimensional space."

To fill in the gaps, Humbird used 80 percent of the Trinity simulations to train a machine-learning model, which was then tested on the remaining 20 percent of the data to evaluate its predictive capability. The model—a random forest decision tree method—accurately predicted yield with a less than 10 percent margin of error. Having a trained machine-learning model in hand that closely mimicked the expensive physics code, Humbird began predicting implosion performance between the simulated data points to search for a robust implosion. "This is not something we could have done with just our physics code," says Humbird, whose work with Peterson has led to a machine-learning project under the Laboratory Directed Research and Development Program. "Performing this search through nine-dimensional space would have required something like 5 million physics simulations and 3 billion central processing unit hours. One would never have enough time. However, a rapid, accurate machine-learning model could do the same search in a fraction of the time."

A sort of topographical map of target designs began to emerge as the model filled in additional data points. Some regions of the map represented locations where an implosion

In this two-dimensional conceptual illustration, varying combinations of target design parameters produce high-yield simulations (light areas) and low-yield simulations (dark areas). The larger the light area, the more resistant the target's implosion would be to perturbations.



Researchers discovered a ridge of high energy yield (yellow) that was larger for an ovoid target than for a sphere. Ignition was found to be far more likely with an ovoid, within a broader range of parameters.

would likely produce high energy yield, whereas other areas indicated the opposite. The high-yield areas indicated how robust the implosion was—what Peterson calls "wobble room." A target designed within a broad "plateau" of high-yield implosions would be resistant enough to withstand perturbations unavoidable in experiments, whereas a target based on a narrow peak on the map might be easily disrupted and "fall off the mountain." After searching through the most promising Trinity simulations and the adjacent machine-learning predictions, the model had what looked like an answer. However, this optimum target did not look like the long-desired sphere but rather more like an egg.

Off Target Can Still Be on Target

The topographical map—much more detailed thanks to Humbird's model—indicated that areas representing egg- or football-shaped targets, known as ovoids, were plateaus of stability. Even if NIF could not create an implosion at the absolute center of the plateau, being slightly off-center would still produce a high energy yield. With a better idea of where to look, Peterson

and Humbird ran a more expensive, full-scale physics simulation on the ovoid target, and their predictions were confirmed, although the researchers were initially not sure why.

After puzzling over the contradictions, Peterson realized he was seeing zonal flows in the imploding egg-shaped target. Similar to a spiraling hurricane sucking up neighboring clouds, zonal flows can absorb disruptions caused by target support tents or capsule roughness and incorporate them into a larger, more stable vortex. This incorporation steadies the implosion and allows for greater energy output, the researchers concluded. The next steps are to improve the detail of future simulations and continue the search for the perfect target shape.

"Our codes indicate that other designs could exist out there, which is fascinating because we've been chasing the same design for 40 years," says Peterson. "The crazy thing is, we didn't force the code to produce the data. The code could always have yielded these results if we had just known where to look. Machine learning and data science gave us the power."

—Ben Kennedy

Key Words: inertial confinement fusion (ICF), Laboratory Directed Research and Development Program, machine learning, National Ignition Facility (NIF), supercomputing, target design, Trinity, zonal flow.

For further information, contact Luc Peterson at (925) 423-5459 (peterson76@llnl.gov).