

Deep learning expands study of nuclear waste remediation

November 13 2019, by Kathy Kincade



A schematic of the physics-informed generative adversarial network used to estimate parameters and quantify uncertainty in the subsurface flow at the Hanford Site. Credit: Prabhat, Lawrence Berkeley National Laboratory

A research collaboration between Lawrence Berkeley National Laboratory (Berkeley Lab), Pacific Northwest National Laboratory (PNNL), Brown University, and NVIDIA has achieved exaflop performance on the Summit supercomputer with a deep learning application used to model subsurface flow in the study of nuclear waste



remediation. Their achievement, which will be presented during the "Deep Learning on Supercomputers" workshop at SC19, demonstrates the promise of physics-informed generative adversarial networks (GANs) for analyzing complex, large-scale science problems.

"In science we know the <u>laws of physics</u> and observation principles—mass, momentum, energy, etc.," said George Karniadakis, professor of applied mathematics at Brown and co-author on the SC19 workshop paper. "The concept of physics-informed GANs is to encode prior information from the physics into the neural network. This allows you to go well beyond the training domain, which is very important in applications where the conditions can change."

GANs have been applied to model human face appearance with remarkable accuracy, noted Prabhat, a co-author on the SC19 paper who leads the Data and Analytics Services team at Berkeley Lab's National Energy Research Scientific Computing Center. "In science, Berkeley Lab has explored the application of vanilla GANs for creating synthetic universes and particle physics experiments; one of the open challenges thus far has been the incorporation of physical constraints into the predictions," he said. "George and his group at Brown have pioneered the approach of incorporating physics into GANs and using them to synthesize data—in this case, subsurface flow fields."

For this study, the researchers focused on the Hanford Site, established in 1943 as part of the Manhattan Project to produce plutonium for nuclear weapons and eventually home to the first full-scale plutonium production reactor in the world, eight other nuclear reactors, and five plutonium-processing complexes. When plutonium production ended in 1989, left behind were tens of millions of gallons of radioactive and chemical waste in large underground tanks and more than 100 square miles of contaminated groundwater resulting from the disposal of an estimated 450 billion gallons of liquids to soil disposal sites. So for the



past 30 years the U.S. Department of Energy has been working with the Environmental Protection Agency and the Washington State Department of Ecology to clean up Hanford, which is located on 580 square miles (nearly 500,000 acres) in south-central Washington, whole parts of it adjacent to the Columbia River—the largest river in the Pacific Northwest and a critical thoroughfare for industry and wildlife.



Using physics-informed GANs on the Summit supercomputer, the research team was able to estimate parameters and quantify uncertainty in subsurface flow. This image shows locations of sensors around the Hanford Site for levels 1 (black) and 2 (color). Units are in km. Credit: Lawrence Berkeley National Laboratory

To track the cleanup effort, workers have relied on drilling wells at the Hanford Site and placing sensors in those wells to collect data about geologic properties and groundwater flow and observe the progression of contaminants. Subsurface environments like the Hanford Site are very



heterogeneous with varying in space properties, explained Alex Tartakovsky, a computational mathematician at PNNL and co-author on the SC19 paper. "Estimating the Hanford Site properties from data only would require more than a million measurements, and in practice we have maybe a thousand. The laws of physics help us compensate for the lack of data."

"The standard parameter estimation approach is to assume that the parameters can take many different forms, and then for each form you have to solve subsurface flow equations perhaps millions of times to determine parameters best fitting the observations," Tartakovsky added. But for this study the research team took a different tack: using a physicsinformed GAN and <u>high performance</u> computing to estimate parameters and quantify uncertainty in the subsurface flow.

For this early validation work, the researchers opted to use synthetic data—data generated by a computed model based on expert knowledge about the Hanford Site. This enabled them to create a virtual representation of the site that they could then manipulate as needed based on the parameters they were interested in measuring—primarily hydraulic conductivity and hydraulic head, both key to modeling the location of the contaminants. Future studies will incorporate real sensor data and real-world conditions.

"The initial purpose of this project was to estimate the accuracy of the methods, so we used synthetic data instead of real measurements," Tartakovsky said. "This allowed us to estimate the performance of the physics-informed GANS as a function of the number of measurements."

In training the GAN on the Summit supercomputer at the Oak Ridge Leadership Computing Facility OLCF, the team was able to achieve 1.2 exaflop peak and sustained performance—the first example of a largescale GAN architecture applied to SPDEs. The geographic extent, spatial



heterogeneity, and multiple correlation length scales of the Hanford Site required training the GAN model to thousands of dimensions, so the team developed a highly optimized implementation that scaled to 27,504 NVIDIA V100 Tensor Core GPUs and 4,584 nodes on Summit with a 93.1% scaling efficiency.

"Achieving such a massive scale and performance required full stack optimization and multiple strategies to extract maximum parallelism," said Mike Houston, who leads the AI Systems team at NVIDIA. "At the chip level, we optimized the structure and design of the neural network to maximize Tensor Core utilization via cuDNN support in TensorFlow. At the node level, we used NCCL and NVLink for high-speed data exchange. And at the system level, we optimized Horovod and MPI not only to combine the data and models but to handle adversary parallel strategies. To maximize utilization of our GPUs, we had to shard the data and then distribute it to align with the parallelization technique."

"This is a new high-water mark for GAN architectures," Prabhat said. "We wanted to create an inexpensive surrogate for a very costly simulation, and what we were able to show here is that a physicsconstrained GAN architecture can produce spatial fields consistent with our knowledge of physics. In addition, this exemplar project brought together experts from subsurface modeling, applied mathematics, deep learning, and HPC. As the DOE considers broader applications of deep learning—and, in particular, GANs—to simulation problems, I expect multiple research teams to be inspired by these results."

The paper, "<u>Highly Scalable, Physics-Informed GANs for Learning</u> <u>Solutions of Stochastic PDEs</u>," will be presented at the SC19 Deep Learning on Supercomputers workshop. I

More information: Highly-scalable, physics-informed GANs for learning solutions of stochastic PDEs, arXiv:1910.13444 [physics.comp-



ph] <u>arxiv.org/abs/1910.13444</u>

Provided by Lawrence Berkeley National Laboratory

Citation: Deep learning expands study of nuclear waste remediation (2019, November 13) retrieved 3 May 2024 from <u>https://phys.org/news/2019-11-deep-nuclear-remediation.html</u>

This document is subject to copyright. Apart from any fair dealing for the purpose of private study or research, no part may be reproduced without the written permission. The content is provided for information purposes only.