

Deep learning stretches up to scientific supercomputers

August 14 2018



Researchers delivered a 15-petaflop deep-learning software and ran it on Cori, a supercomputer at the National Energy Research Scientific Computing Center, a Department of Energy Office of Science user facility. Credit: Lawrence Berkeley National Laboratory

Machine learning, a form of artificial intelligence, enjoys unprecedented success in commercial applications. However, the use of machine learning in high performance computing for science has been limited. Why? Advanced machine learning tools weren't designed for big data



sets, like those used to study stars and planets. A team from Intel, National Energy Research Scientific Computing Center (NERSC), and Stanford changed that situation. They developed the first 15-petaflop deep-learning software. They demonstrated its ability to handle large data sets via test runs on the Cori supercomputer.

Using machine learning techniques on supercomputers, scientists could extract insights from large, complex <u>data sets</u>. Powerful instruments, such as accelerators, produce massive data sets. The new software could make the world's largest supercomputers able to fit such data into deep learning uses. The resulting insights could benefit Earth systems modeling, fusion energy, and astrophysics.

Machine learning techniques hold potential for enabling scientists to extract valuable insights from large, complex data sets being produced by accelerators, light sources, telescopes, and computer simulations. While these techniques have had great success in a variety of <u>commercial applications</u>, their use in <u>high performance computing</u> for science has been limited because existing tools were not designed to work with the terabyte- to petabyte-sized data sets found in many science domains.

To address this problem a collaboration among Intel, the National Energy Research Scientific Computing Center, and Stanford University has been working to solve problems that arise when using deep learning techniques, a form of machine learning, on terabyte and petabyte data sets. The team developed the first 15-petaflop deep-learning software. They demonstrated its scalability for data-intensive applications by executing a number of training runs using large scientific data sets. The runs used physics- and climate-based data sets on Cori, a supercomputer located at the National Energy Research Scientific Computing Center. They achieved a peak rate between 11.73 and 15.07 petaflops (singleprecision) and an average sustained performance of 11.41 to 13.47



petaflops. (A petaflop is million billion calculations per second.)

More information: Thorsten Kurth et al. Deep learning at 15PF, *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis on - SC '17* (2017). DOI: 10.1145/3126908.3126916

Provided by US Department of Energy

Citation: Deep learning stretches up to scientific supercomputers (2018, August 14) retrieved 26 April 2024 from <u>https://phys.org/news/2018-08-deep-scientific-supercomputers.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.