

Speed Up Your Machine Learning

GPU-enabled systems accelerate machine learning & deep learning

We simplify our world and speed up problem-solving using heuristics — shortcuts that help us make decisions based on experience or similarity. Machine learning and deep learning use heuristic programming, where observations and data analysis are converted into patterns and rules for easier processing.

To effectively use heuristics, your machine learning and deep learning problems need technology that makes quick work of GPU-enabled applications and gives you the highest performance per dollar and per watt. Reducing training times for deep learning models can make a critical difference in your results.

Cray and NVIDIA Make It Happen

Cray delivers a powerful deep learning package in a single system, with NVIDIA GPUs for machine learning, deep learning and extreme HPC. You can even use different Cray platforms together to enable complete multi-stage workflows. We've worked with leading deep learning packages on systems that scale beyond typical GPU density.



GPU-ACCELERATED COMPUTING

The Cray® CS-Storm™ cluster supercomputer is a flexible, open and integrated system for machine learning and extreme HPC computing, and the foundation for organizations that want to take advantage of GPU accelerators in their production workflows. Our CS-Storm series includes support for up to 10 NVIDIA® Tesla® P40 or P100 PCIe accelerators, or eight Tesla P100 SXM2 accelerators utilizing the NVIDIA® NVLink™ high speed interconnect.



SCALABLE DEEP LEARNING

The Cray® XC50™ accelerated GPU supercomputer combines NVIDIA's latest Tesla P100 PCIe GPUs with Cray's Aries™ network and high-performance software environment. It's our fastest system with peak performance of 500 PF, and 1 PF in a single cabinet. Cray's XC50 system is the world's most scalable supercomputer for deep neural network training.



WE'VE GONE BIG

Cray's deep learning group and engineers at the Swiss National Supercomputing Centre (CSCS) successfully scaled out the Microsoft Cognitive Toolkit — an open-source, commercial-grade toolkit that trains deep learning algorithms to learn like the human brain — to over 1,024 Cray XC50 supercomputer nodes.

Scaling out deep learning

“Cray's proficiency in performance analysis and profiling, combined with the unique architecture of the XC systems, allowed us to bring deep learning problems to our Piz Daint system and scale them in a way that nobody else has. What is most exciting is that our researchers and scientists will now be able to use our existing Cray XC supercomputer to take on a new class of deep learning problems that were previously infeasible.”

- Prof. Dr. Thomas C. Schulthess, director of the Swiss National Supercomputing Centre (CSCS)