The Laboratory's New AI Supercomputer is the Most Powerful at any University

Division 5: Cyber Security and Information Sciences

The new TX-GAIA computing system at the Lincoln Laboratory Supercomputing Center (LLSC) has been ranked as the most powerful artificial intelligence (AI) supercomputer at any university in the world. The ranking comes from TOP500, which publishes a list of the top supercomputers in various categories biannually. The system, which was built by Hewlett Packard Enterprises, combines traditional high-performance computing hardware (nearly 900 Intel processors) with hardware optimized for AI applications (900 Nvidia GPU accelerators).

"We are thrilled by the opportunity to enable researchers across the Laboratory and MIT to achieve incredible scientific and engineering breakthroughs," said Dr. Jeremy Kepner, Laboratory Fellow, LLSC. "TX-GAIA will play a large role in supporting AI, physical simulation, and data analysis across all Laboratory missions."

TOP500 rankings are based on the LINPACK Benchmark, which is a measure of a system's floating-point computing power, or how fast a computer solves a dense system of linear equations. TX-GAIA's TOP500 benchmark performance of 3.9 quadrillion floating-point operations

TX-GAIA is housed inside of a new EcoPOD deployable data center, manufactured by Hewlett Packard Enterprises, at the site of the Lincoln Laboratory Supercomputing Center in Holyoke, Massachusetts.

per second, or petaflops, places it #1 in the Northeast, #20 in the United States, and #51 in the world for supercomputing power. The system's peak performance is more than 6 petaflops.

But more notably, TX-GAIA has a peak performance of 100 AI petaflops, which makes it #1 for AI flops at any university in the world. An AI flop is a measure of how fast a computer can perform deep neural network (DNN) operations. DNNs are a class of AI algorithms that learn to recognize patterns in huge amounts of data. This ability has given rise to "AI miracles," as Kepner put it, in speech recognition and computer vision; the technology is what allows Amazon's Alexa to understand questions and self-driving cars to recognize objects in their surroundings. As these DNNs grow more complex, it takes longer

for them to process the massive datasets they learn from. TX-GAIA's Nvidia GPU accelerators are specially designed for performing these DNN operations quickly.

TX-GAIA is housed in a new modular data center, called an EcoPOD, at the site of the LLSC in Holyoke, Massachusetts. It joins the ranks of other powerful systems at the LLSC, such as the TX-E1, which supports collaborations with MIT campus and other institutions, and TX-Green, which is currently ranked #492 on the TOP500 list and is a zero-carbon system.

Kepner said that the system's integration into the LLSC will be completely transparent to users when it comes online this fall: "The only thing users should see is that many of their computations will be dramatically faster."

Among its AI applications, TX-GAIA will be tapped for training machine-learning algorithms, including those that use DNNs. It will more quickly crunch through terabytes of data—for example, hundreds of thousands of images or years' worth of speech samples to teach these algorithms to figure out

"The only thing users should see is that many of their computations will be dramatically faster."

solutions on their own. The system's compute power will also expedite simulations and data analysis. These capabilities will support projects across the Laboratory's R&D areas, such as improving weather forecasting, accelerating medical data analysis, building autonomous

systems, designing synthetic DNA, and developing new materials and devices.

TX-GAIA, which is also ranked the #1 system in the Department of Defense, will also support the recently announced MIT-Air Force AI Accelerator. The partnership will

combine the expertise and resources of MIT, including those at the LLSC and the Air Force, to conduct fundamental research directed at enabling rapid prototyping, scaling, and application of AI algorithms and systems.

MIT LINCOLN LABORATORY
SUPERCOMPUTING CENTER