

Researchers Deploy GPUs to Build World's Largest Artificial Neural Network

GPU-Accelerated Machine Learning and Data Mining Poised to Dramatically Improve Object, Speech, Audio, Image and Video Recognition Capabilities

LEIPZIG, GERMANY -- ISC 2013 -- NVIDIA today announced that it has collaborated with a research team at Stanford University to create the world's largest artificial neural network built to model how the human brain learns. The network is 6.5 times bigger than the previous record-setting network developed by Google in 2012.

Computer-based neural networks are capable of "learning" how to model the behavior of the brain -- including recognizing objects, characters, voices and audio in the same way that humans do.

Yet creating large-scale neural networks is extremely computationally expensive. For example, Google used approximately 1,000 CPU-based servers, or 16,000 CPU cores, to develop its neural network, [which taught itself to recognize cats](#) in a series of YouTube videos. The network included 1.7 billion parameters, the virtual representation of connections between neurons.

In contrast, the Stanford team, led by Andrew Ng, director of the university's Artificial Intelligence Lab, created an equally large network with only three servers using NVIDIA® GPUs to accelerate the processing of the big data generated by the network. With 16 NVIDIA GPU-accelerated servers, the team then created an 11.2 billion-parameter neural network -- 6.5 times bigger than a network Google announced in 2012.

The bigger and more powerful the neural network, the more accurate it is likely to be in tasks such as object recognition, enabling computers to model more human-like behavior. A paper on the Stanford research was published yesterday at the [International Conference on Machine Learning](#).

"Delivering significantly higher levels of computational performance than CPUs, GPU accelerators bring large-scale neural network modeling to the masses," said Sumit Gupta, general manager of the Tesla Accelerated Computing Business Unit at NVIDIA. "Any researcher or company can now use machine learning to solve all kinds of real-life problems with just a few GPU-accelerated servers."

GPU Accelerators Power Machine Learning

[Machine learning](#), a fast-growing branch of the artificial intelligence (AI) field, is the science of getting computers to act without being explicitly programmed. In the past decade, machine learning has given us self-driving cars, effective web search and a vastly improved understanding of the human genome. Many researchers believe that it is the best way to make progress towards human-level AI.

One of the companies using GPUs in this area is Nuance, a leader in the development of speech recognition and natural language technologies. Nuance trains its neural network models to understand users' speech by using terabytes of audio data. Once the models are trained, they can then recognize the pattern of spoken words by relating them to the patterns that the model learned earlier.

"GPUs significantly accelerate the training of our neural networks on very large amounts of data, allowing us to rapidly explore novel algorithms and training techniques," said Vlad Sejnoha, chief technology officer at Nuance. "The resulting models improve accuracy across all of Nuance's core technologies in healthcare, enterprise and mobile-consumer markets."

NVIDIA will be exhibiting at the [2013 International Supercomputing Conference \(ISC\)](#) in Leipzig, Germany this week, June 16-20, at booth #220.

About NVIDIA

Since 1993, [NVIDIA](#) (NASDAQ: NVDA) has pioneered the art and science of visual computing. The company's technologies are transforming a world of displays into a world of interactive discovery -- for everyone from gamers to scientists, and consumers to enterprise customers. More information at <http://nvidianews.nvidia.com> and <http://blogs.nvidia.com>.

Certain statements in this press release including, but not limited to, statements as to: the impact and benefits of NVIDIA GPU accelerators are forward-looking statements that are subject to risks and uncertainties that could cause results to be materially different than expectations. Important factors that could cause actual results to differ materially include: global economic conditions; our reliance on third parties to manufacture, assemble, package and test our products; the impact of technological development and competition; development of new products and technologies or enhancements to our existing product and technologies; market acceptance of our products or our partners' products; design, manufacturing or software defects; changes in consumer preferences or demands; changes in industry standards and interfaces; unexpected loss of performance of our products or technologies when integrated into systems; as well as other factors detailed from time to time in the reports NVIDIA files with the Securities and Exchange Commission, or SEC, including its Form 10-Q for the fiscal period ended April 28, 2013. Copies of reports filed with the SEC are posted on the company's website and are available from NVIDIA without charge. These forward-looking statements are not guarantees of future performance and speak only as of the date hereof, and, except as required by law, NVIDIA disclaims any obligation to update these forward-looking statements to reflect future events or circumstances.

© 2013 NVIDIA Corporation. All rights reserved. NVIDIA and the NVIDIA logo are trademarks and/or registered trademarks of NVIDIA Corporation in the U.S. and other countries. Other company and product names may be trademarks of the respective companies with which they are associated. Features, pricing, availability and specifications are subject to change without notice.

About NVIDIA

Since 1993, [NVIDIA](#) (NASDAQ: NVDA) has pioneered the art and science of [visual computing](#). The company's technologies are transforming a world of displays into a world of interactive discovery — for everyone from gamers to scientists, and consumers to enterprise customers. More information at <http://nvidianews.nvidia.com/> and <http://blogs.nvidia.com/>.

© 2014 NVIDIA Corporation. All rights reserved. NVIDIA and the NVIDIA logo are trademarks and/or registered trademarks of NVIDIA Corporation in the U.S. and other countries. Other company and product names may be trademarks of the respective companies with which they are associated. Features, pricing, availability, and specifications are subject to change without notice.

Media Contacts

Hector Martinez

+1 408 486 3443

hmarinez@nvidia.com

George Millington

+1 408 562 7226

gmillington@nvidia.com