

7 Ways GPUs are Accelerating Data Science

Artificial intelligence (AI) is unlocking new frontiers for the academic and scientific communities and the world – and is doing so faster than ever before.

AI was previously constrained by the limitations of slow compute architectures based on central processing units (CPUs), which cannot meet the demands of today's massively parallelised deep-learning algorithms. Computer systems taking advantage of the chip architecture of graphics processing units (GPUs) are changing all that.

While a CPU is composed of just a few cores with lots of cache memory that can handle a few software threads at a time, a GPU is composed of hundreds of cores that can handle thousands of threads simultaneously. This enables GPUs to

break complex problems into thousands or millions of separate tasks and work them out at once.

NVIDIA originally invented the GPU to render graphics through rapid mathematical calculations, and it's that high-performance processing that makes them such powerful workhorses for data science.

Through the power of GPU-accelerated deep learning, students and researchers can evolve and embrace data science to fully realise the potential of AI across numerous disciplines.

This is just a taste of the kind of data-science projects that GPUs are supporting at universities, research institutes and not-for-profit organisations around the world.

TURBINE OR NOT TURBINE – ASSESSING WIND TURBULENCE AT DEAKIN UNIVERSITY

How much of a turbulence wake do wind turbines create, and how can this be measured in relation to the hazard they may pose to aviation? Jorg Schluter from Deakin University in Melbourne and Sindhu Paramasivam from Nanyang Technological University in Singapore set out to find out, in collaboration with the Australian Civil Aviation Safety Authority (CASA) and wind-farm developers.

The team used Large Eddy Simulations (LES) to model the turbulence of wind-turbine wakes to better quantify the aviation hazard that these wakes and turbulence might cause to small aircraft. GPU computing played a key role in processing the large datasets and enabling real-time visualisations.

A typical simulation uses 50 million mesh points, saved over 2,000 time steps, creating over 1TB of data for each run. Running LES models takes large amounts of computing resources and the visualisation of these data sets – in real time for engineering analysis or to create animations – has traditionally been very costly and slow, which has impeded investigations and accurate simulations. By deploying GPU resources, the team was able to accelerate the rendering of the 3D data sets by 30% and successfully visualise the turbulence streams, enabling them to analyse the potential hazard to small planes and better inform wind-farm planning.

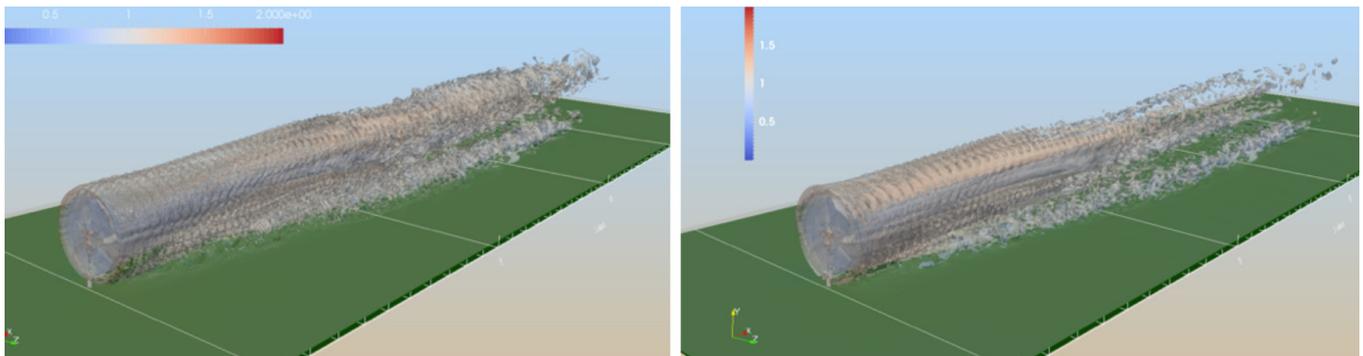


Figure 3: *Isosurfaces of vorticity of the NREL 5MW turbine. Left: LES solution; Right: LES solution filtered with a Gaussian filter of 10m filter width.*

Image from Hazard Assessment of Wind Turbine Wakes Turbulence: Initial Results, Schulte, Jorg and Paramasivam, Sindhu, 2018

BOOSTING NEURAL-NETWORK TRAINING BY 80% AT DARTMOUTH COLLEGE

A team of researchers at Dartmouth College in New Hampshire, USA, upgraded to a new NVIDIA GPU and unlocked massive performance improvements. The Hassanpour Lab at the college is focused on ‘harnessing the power of data for precision health’. Researchers design and develop methods to assist radiologists and pathologists, such as deep learning-based image processing systems to identify subtle findings in

large volumes of radiology images or the analysis of high-resolution microscopy images. Training deep-learning models requires powerful hardware resources, and every new iteration of NVIDIA GPUs improves their outcomes. Running their existing code on the new GPU, the team achieved an 80% performance increase when training a pair of neural networks to detect osteoporotic vertebral fractures.



DARTMOUTH

VITAL LIFE SCIENCES RESEARCH ACCELERATED BY 50% AT CSL

Biotech giant CSL turned to GPUs when they contracted XENON for a HPC solution to support their world-leading research. The brief from CSL was to improve the speed and capabilities of research projects, eliminate processing bottlenecks even as datasets grew, enable faster data analytics for projects such as drug trials, and to build a long-term technology platform to accommodate the increasing demand from burgeoning genomics and biotech data.

XENON's solution was a new HPC cluster integrated with the existing environment, and included a GPU server loaded with four NVIDIA V100 GPUs. The increased compute and processing power quickly opened new research opportunities for CSL scientists.

CSL's IT system holds a vast array of research tools, spanning the basics of data analytics and software development, as well as sophisticated image-processing tools and applications for genomics. The new platform features GPU accelerators to speed up applications involving advanced analytics techniques such as machine learning, deep learning and AI, allowing hundreds of research projects to run simultaneously. The upgrade vastly reduced the analytics timeline, with some projects running 50% faster on the new platform. Further, when users harness multiple nodes of the cluster, time-frames shrink from days to hours – something not previously possible.



MODELLING ARTIFICIAL EYES AT THE BIONIC VISION LAB AT UC SANTA BARBARA

More than 10 million people around the world live with profound visual impairment. Retinal neuroprostheses are being developed at the Bionic Vision Lab at UC Santa Barbara to restore their vision. Prosthetic vision, however, is still rudimentary, so rather than aiming to restore natural vision, the Bionic Vision Lab is using deep learning-based scene simplification, borrowing

state-of-the-art computer vision algorithms as image-processing techniques to maximise the usefulness of prosthetic vision. The Bionic Vision Lab uses deep learning and virtual prototyping on NVIDIA GPUs to develop models of artificial eyes, and researchers explore the potential and limits of a design for artificial eyes by viewing a model through a virtual-reality headset.



PROTECTING THE OCEANS WITH DRONES AND GPUS AT ATLAN SPACE

In Africa, the ATLAN Space project uses a fleet of autonomous drones with AI-powered computer vision to detect illegal fishing and ships dumping oil into the sea. ATLAN Space is a member of NVIDIA's Inception virtual accelerator program and uses NVIDIA GPUs to train its neural networks and the NVIDIA Jetson embedded platform for inference.

model spots a boat, it analyses the image to identify whether it's a fishing vessel. Next, the neural network analyses the boat's name, flag and type of radio signals to determine whether it is legally permitted to operate in the region. All this inference work happens on Jetson, processing even if the drone loses its satellite connection during flight.



A single drone allows them to monitor 10,000 square kilometres a day, while navigating at a closer range, under any clouds. ATLAN Space deploys its AI on a Jetson TX2 board, which connects to the drone's autopilot and communication system. The autonomous drone flies on a path determined by the neural network as it scans the seas. Once the deep-learning

When the neural network identifies a probable unauthorised boat, it alerts authorities via a satellite message. Back on land, the raw data collected by the drones is fed into software running on NVIDIA GPUs through the Microsoft Azure cloud platform. The deep neural networks learn from this additional data to improve and optimise for future missions.

ACCELERATING AUTONOMOUS VEHICLE RESEARCH AT CLEMSON UNIVERSITY

A group of researchers at Clemson University in South Carolina, USA is working on the Open Connected Automated Vehicle (OpenCAV) project. The Metamoto simulation software they're using to build an augmented reality (AR) component requires GPU acceleration to deliver optimal performance. Clemson University has one of the largest public academic supercomputers in the US, the Palmetto cluster. It's been built over more than a decade and has more than 2000 compute nodes and more than 1000 NVIDIA GPUs. Clemson's IT administrators pride

themselves on “zero red tape”, with all students and faculty researchers able to access Palmetto for whatever job they want. For the OpenCAV researchers, they found the value of NVIDIA's virtual GPU (vGPU) technology, which allowed them to deliver resources for specialised workload characteristics and resource requirements. The OpenCAV team uses NVIDIA V100 GPUs and NVIDIA RTX Virtual Workstation (vWS) software, which lets researchers do their work faster with one-quarter of the physical hardware that would otherwise have been required.



SPEEDING UP COMPUTATIONS AT CSIRO

More than a decade ago, XENON delivered Australia's first GPU HPC cluster for CSIRO. Three subsequent NVIDIA GPU upgrades made the BRAGG, in Canberra, then the 10th most energy-efficient supercomputer in the world. XENON rolled out progressive upgrades to increase

capacity and minimise environmental impact and energy costs.

The productivity boost launched as soon as it was installed: The GPU HPC cluster allowed CSIRO scientists to perform computations in a single morning that used to take weeks.



XENON – ENABLING DATA SCIENTISTS TO DO GREAT NEW THINGS

XENON delivers for data science teams looking to iterate faster, collaborate more, and develop solutions for tomorrow's problems today. As XENON has for over 25 years.

Contact XENON:

- 1300 888 030 (Australia)
- +61 3 9549 1111 (International)
- info@xenon.com.au
- www.xenon.com.au

Explore the resources below, or get in touch with the XENON team to explore how NVIDIA GPU technology can assist your team.

[Talk to a Solutions Architect](#)

REFERENCES AND RELATED READING

- **Atlan Space**
<http://www.atlanspace.com/>
- **Bionic Vision Lab**
<https://bionicvisionlab.org/#research>
- **Clemson University Autonomous Vehicle research**
<https://www.youtube.com/watch?v=uACVoNThBdk&list=PLZHnYvH1qtOYZsJekEqxWV9pQ7lIBHUJF&index=2>
- **CSIRO and XENON case studies**
<https://xenon.com.au/tag/csiro/>
- **CSL Accelerating Research Outcomes**
<https://xenon.com.au/case-studies/xenon-and-csl-accelerating-research-outcomes/>
- **NVIDIA Inception Program**
<https://www.nvidia.com/en-au/deep-learning-ai/startups/>