



Data Drives Innovation Even Further When It's Backed By GPUs

Modern research runs on data and it's now the norm that massive amounts of it are collected from every corner of daily life. Data scientists explore these vast tranches of digital information to extract insights and to build, train and deploy predictive models using artificial intelligence (AI) and machine learning (ML).

The world is changing because of it. Data scientists work across every discipline and purpose, from better understanding customer behaviour in order to assist a retail enterprise to ground-breaking discoveries in medical science that improve the diagnosis and treatment of life-threatening diseases.

This level of data science requires massive compute power, and graphics processing units (GPUs) are today the accelerator of choice for scientific researchers building deep learning and AI models. With these workloads, the central processing unit (CPU) is still essential, but has been relieved of most intensive computational tasks that require parallel processing.

For 30 years, the dynamics of Moore's Law held true as microprocessor performance grew at 50% per year. But the limits of semiconductor physics mean that CPU performance now only grows by 10% per year. Back when data scientists relied on CPUs to load, filter and manipulate data and train and deploy models, workflows were slow and cumbersome and could take many days just for a single run, especially with large, complex models involving a lot of features.

NVIDIA invented the GPU in 1999 and its expertise in programmable GPUs has since led to breakthroughs in parallel processing which have made supercomputing widely accessible. NVIDIA GPU computing gives data science a path forward and will provide a 1,000X speed-up by 2025 when compared to CPU processing of the same workloads.

The advantage of the NVIDIA suite of GPUs is that a relatively simple upgrade with one of them gives an existing system an instant and enormous boost to data-processing power, transforming the work of users.



This GPU technology today underpins data science brilliance around the world, allowing practitioners across all disciplines to push their ideas to new limits at unprecedented speeds.

GPUs reduce infrastructure costs and they provide superior performance for end-to-end data science workflows when using RAPIDS™ open-source software libraries. GPU-accelerated data science is available everywhere – on the laptop, in the data center, and in the cloud.

NVIDIA focuses on accelerating data science by offering the only integrated hardware-to-software stack that comes already optimised for data science, meaning that it offers vastly improved performance as soon as it's installed, and the invaluable insights gained from faster processing will be delivered much, much sooner.

How much sooner? GPUs and RAPIDS can accelerate ML training by up to 215X faster than CPU alone, performing more iterations, opening up the opportunity for more experimentation and deeper exploration.

At the same time, infrastructure costs are reduced and data-centre efficiency is increased. Data scientists, students and enterprises using on-premise GPUs don't have to count how many hours of system use they're racking up or budget how many runs they can afford over a particular timespan.

If a new methodology fails at first, there's no added investment required to try a different variation of code, encouraging experimentation and creativity. In fact, the more the GPU-based system is used for this

work, the greater the return on investment. From powerful desktop GPUs to workstations and enterprise systems, NVIDIA's GPU solutions for data science come in a broad spectrum of choices to meet specific workloads and operational requirements. Depending on the price and performance needs, developers might start off with a single NVIDIA GPU in a workstation and eventually ramp up to a cluster of GPU based servers. With each GPU upgrade, work can continue seamlessly.

The growth of applications that are being accelerated by NVIDIA GPUs has created a powerful ecosystem. NVIDIA has created a catalogue of GPU-accelerated applications, which span every academic discipline and all industries. Here are just a few.

HEALTHCARE AND LIFE SCIENCES

Healthcare demands new computing paradigms to meet the need for personalised medicine, next-generation clinics, enhanced quality of care, and breakthroughs in biomedical research to treat disease. AI healthcare startups are vital to pushing the boundaries of health innovation. NVIDIA Inception nurtures more than 1,000 healthcare startups developing cutting-edge, GPU-based tools to optimise operations, enhance diagnostics, and elevate patient care. NVIDIA Clara™ is a healthcare application framework for AI-powered imaging, genomics, and the development and deployment of smart sensors and AI-enabled medical devices. It includes a full stack of GPU-accelerated libraries, SDKs, and reference applications for developers, data scientists, and researchers to create real-time, secure, and scalable solutions.



ARCHITECTURE, ENGINEERING AND CONSTRUCTION (AEC)

The design, construction, and operation of buildings and infrastructure brings complex challenges. To meet them, architecture, engineering, and construction (AEC) companies worldwide use NVIDIA technologies to optimize designs, mitigate hazards, and collaborate more effectively, even when working remotely. With breakthroughs in AI, 3D graphics virtualisation, virtual reality (VR), and collaboration solutions such as NVIDIA Omniverse™, firms are transforming their workflows. With GPU-powered data science and AI, AEC teams tap into vast amounts of data – including design, environmental, simulation, and structural – to gain insights into everything from traffic flows to weather conditions.

HIGHER EDUCATION AND RESEARCH

Higher education and research are at the forefront of solving the grand challenges of our time leveraging the tools of AI, accelerated computing, and data science. Institutions must also meet the demand for an AI-fluent workforce. This entails not only preparing future generations of workers, but also creating challenging new curriculum based on fast moving

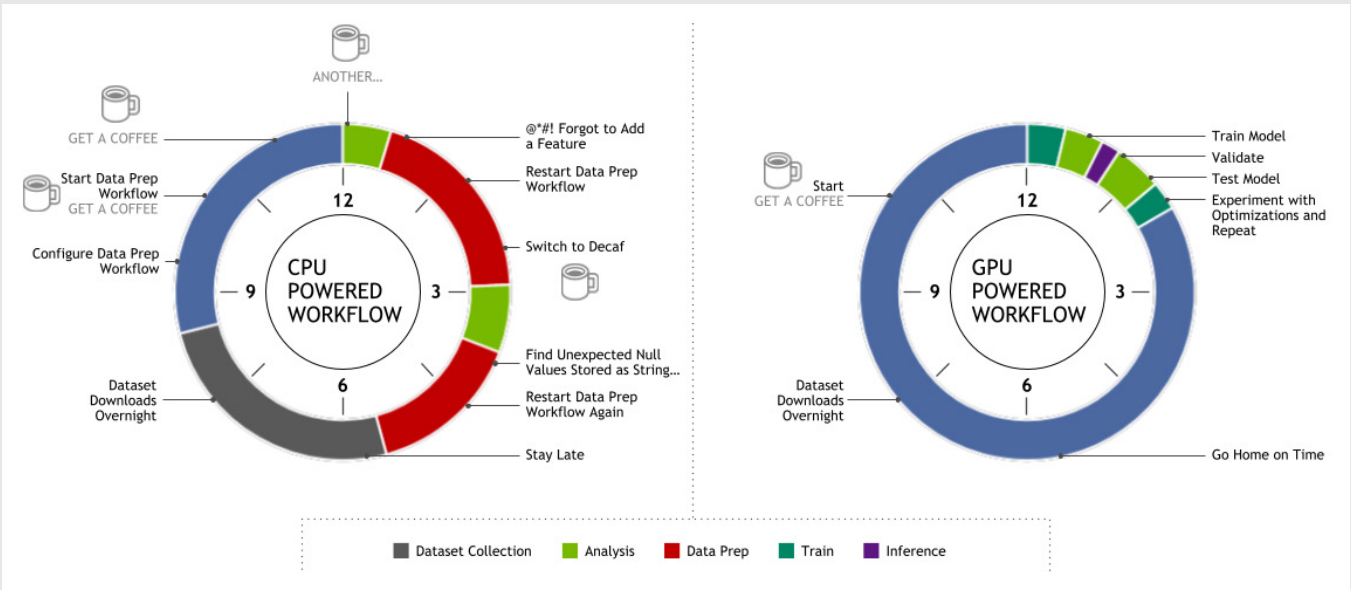
research using new teaching frameworks and tools. From on-premise to the cloud, NVIDIA supplies the tools that accelerate discovery and workforce development wherever they are needed. GPU-accelerated AI and HPC enables researchers to reduce the time to discovery through faster modeling, simulation, and processing of experimental data to tackle even the grandest of challenges.

MANUFACTURING

Data science is reinventing manufacturing. Designers, engineers, and analysts require processing of massive models so they can innovate, iterate, and solve problems – all at the speed of light. Meanwhile, industrial companies need unprecedented amounts of sensor and operational data to optimise operations, improve time to insight, and reduce costs. NVIDIA GPU-accelerated technology makes it all possible, at scale, and from anywhere. Mechanical engineers and engineering analysts leverage the power of NVIDIA's full-stack computing platform to work with massive datasets and perform complex simulations to tackle challenges that were previously impossible. From computer-aided design (CAD) to structural analysis to computational fluid dynamics (CFD), NVIDIA's ecosystem of GPU-accelerated software is fundamentally changing the way manufacturing operates.

A DAY IN THE LIFE OF A DATA SCIENTIST

On the left is a traditional CPU-bound data science workflow: there's a lot of time spent on data prep work or waiting on data prep processes to complete on CPU – denoted by the red. The green portions are actual work done by the data scientist, requiring them to make adjustment, tune, and optimise. There's a lot of downtime involved waiting or doing things that have nothing to do with data science. When a workflow is GPU-accelerated, the vast majority of wait time is eliminated thanks to a dramatically faster data ingest, manipulation and feature engineering that responds immediately to the data scientist and gives them this ability to interact with their datasets in real time. Much more time is spent analysing information and the training itself is also dramatically faster. Each iteration leads to a better, more accurate model – and they're spending way less time drinking coffee as they wait on the machine, and more time using their brilliant brains to build something important.



XENON – ENABLING DATA SCIENTISTS TO DO GREAT NEW THINGS

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

Contact XENON:

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

Explore the resources below, and 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

- **NVIDIA Inception Program**
<https://www.nvidia.com/en-au/deep-learning-ai/startups/>
- **NVIDIA Omniverse**
<https://developer.nvidia.com/nvidia-omniverse-platform>
- **XENON Case Studies**
<https://xenon.com.au/category/case-studies/>
- **A Day in the Life of a Data Scientist**
https://docs.google.com/presentation/d/1bthrPs3kWAZs6_3rbLWs8o74M9lkfyK/edit#slide=id.p6