



Transforming Data Science With Accelerated Computing

NVIDIA RTX PRO AI workstations for data science workflows.



Challenges

Every industry, from retail to financial services to healthcare, is being transformed with AI to improve competitiveness and operational efficiency. In retail, chatbots are simplifying product inventory search for users. In the financial services industry, AI-based tools are enhancing fraud detection. In healthcare, AI is helping uncover new insights from genomics that are leading to better patient outcomes.

AI requires rigorous training on domain-specific datasets to achieve its full potential in production use cases. To prepare those training datasets and production pipelines, data scientists preprocess datasets into tables with millions of rows and columns. Data scientists do this by cleaning the data, managing missing values, converting data types, and engineering specific features needed for the model. Data preprocessing is especially important in domain-specific areas, as training on specific and often proprietary data informs how intelligent the AI becomes in their area.

The good news is that we have more data than ever, with data volume expected to grow at a CAGR of 24.4% from 2023–2028.¹ The problem arising is that traditional CPU-based data processing tools were not designed to handle the scale and complexity of computations required to effectively process today's data volumes and demand. Traditional infrastructure is slowing down, reducing productivity and extending time to market for enterprises.

Accelerate Every Step of Your Data Science Workflow

To meet the demands of preprocessing data for AI, NVIDIA offers a comprehensive full-stack solution that combines accelerated data processing libraries built on NVIDIA® CUDA® with NVIDIA RTX PRO™ AI workstations. CUDA-X™ Data Science is a collection of open-source libraries that accelerate popular data science libraries and platforms, including zero-code-change APIs that accelerate popular PyData tools like pandas, scikit-learn, NetworkX, and XGBoost. This full-stack solution is designed to provide unmatched performance for data science tasks, leveraging the power of parallel processing in NVIDIA RTX PRO Blackwell Generation GPUs. With CUDA-X Data Science on RTX PRO AI workstations, data scientists can use the tools they love in a local environment, making processing bottlenecks a thing of the past.

Key Challenges for Data Science


- **Data Preparation:** Data preparation is a complex, time-consuming process that takes most of a data scientist's time.
- **Scaling:** Volume of data is growing at a rapid pace. Data scientists may resort to downsampling datasets to make large datasets more manageable, leading to suboptimal results.
- **Hardware:** Demand for accelerated AI hardware for data centers and cloud service providers (CSPs) is exceeding supply. Current desktop computing resources may not be suitable for data science workflows..

Benefits of RTX-Powered AI Workstations for Data Science

- Deliver unparalleled acceleration through physical workstations and RTX™ Virtual Workstations.
- Large GPU memory and up to four GPUs per system configuration enable data exploration, feature and model evaluation, and visualization.

1. Source: IDC, Worldwide IDC Global DataSphere Forecast, 2024–2028: AI Everywhere, But Upsurge in Data Will Take Time, Doc #US52076424, May 2024

End-to-End Accelerated Analytics with NVIDIA



- > **Data Preparation:** Rapidly process massive datasets with zero code changes with NVIDIA RTX PRO GPUs and libraries like NVIDIA cuDF and cuML. NVIDIA cuDF accelerates pandas operations by up to 50x performance gains with zero code changes for quick data preparation. This allows you to bypass the limitations of traditional tools and immediately benefit from enhanced processing speeds.
- > **Exploratory Data Analysis:** Extract actionable insights with advanced analytics and visualization tools. CUDA-X Data Science libraries work with popular PyData tools that allow you to create expansive, detailed visualizations that are interactive and responsive to updates and changes. These visualizations enhance data understanding and support critical decision-making.
- > **Model Training:** With GPU-accelerated XGBoost, model training times dramatically reduce from weeks to minutes. This acceleration enables faster iterations, improving operational efficiency and productivity.

High-Performance AI Workstations for Data Science


NVIDIA RTX PRO AI workstations are built to drive high-performance data science. By combining the latest NVIDIA RTX PRO 6000 Blackwell Workstation Edition Series GPUs with a robust software stack—including NVIDIA CUDA-X and data processing and machine learning libraries—these workstations deliver accelerated workflows for data preparation, model training, and visualization.

AI workstations are optimized for fast, local processing with minimal latency, offering a seamless development experience. With NVIDIA AI Workbench, they streamline project setup, development, and collaboration across various environments—from desktops to cloud platforms.

- > Provides additional AI computing resources to augment data center and cloud instances for development and R&D tasks.
- > AI workstations are optimized to leverage NVIDIA CUDA-X Data Science, a collection of open-source libraries that accelerate popular data science libraries and platforms.
- > Enterprise-grade hardware maximizes uptime with enterprise-level performance, reliability, and support.


100x Acceleration

- > **Zero Code Change Acceleration:** Users can enable GPU acceleration by simply loading a magic command in a Jupyter Notebook or using a command-line flag in a Python script.
- > **Unified CPU/GPU Workflows:** Operations are first attempted on the GPU. If the GPU cannot handle the operation, it falls back to the CPU, ensuring compatibility and robustness.
- > **Third-Party Library Compatibility:** CUDA-X data science libraries are designed to be compatible with most third-party libraries in the PyData ecosystem, allowing for seamless integration into existing workflows.
- > **Scalable Development Environment:** With AI workstations, data science workloads can be accelerated from laptop to cloud across teams, enabling seamless collaboration.



CUDA-X Data Science Libraries


For data scientists using pandas, NVIDIA’s cuDF provides a zero-code-change solution that transforms pandas operations, offering up to 50x faster performance on the NVIDIA RTX PRO 6000 Blackwell Workstation Edition. NVIDIA’s accelerated end-to-end data analytics toolchain maximizes data scientist productivity compared to CPU-based solutions.



Specs: Duck DB cudf.Pandas 25.02 on AMD Ryzen Threadripper PRO 7995WX | NVIDIA RTX PRO 6000 Blackwell Workstation Edition | NVIDIA RTX PRO 6000 Blackwell Max-Q Workstation Edition

NVIDIA cuML accelerates machine learning workloads with GPU-optimized implementations of widely used algorithms, delivering massive performance gains with minimal code changes. Designed with a scikit-learn-compatible API, cuML enables seamless migration from CPU-based pipelines while supporting scalable execution of classical ML techniques such as linear regression, clustering, and dimensionality reduction. On the RTX PRO 6000 Blackwell Workstation Edition,


cuML achieves up to 65x faster performance on scikit-learn workloads, 76x faster on HDBSCAN clustering, and 220x faster on UMAP for dimensionality reduction—unlocking rapid experimentation and iteration at enterprise scale.



Specs: NVIDIA cuML 25.02 on NVIDIA RTX PRO 6000 Blackwell Workstation Edition | scikit-learn 1.5 on AMD RYZEN



Desktop	Good	Better	Best
CPU	Intel Core Ultra 9 Series	AMD Threadripper Pro or Intel Xeon w5	AMD Threadripper Pro or Intel Xeon w9
System Memory	192 GB ECC DDR5	384 GB ECC DDR5 (or more)	758 GB ECC DDR5 (or more)
Storage	1 TB boot + 2 TB NVMe SSDs	1 TB boot + 2--4 TB NVMe SSDs with RAID ²	2 TB boot + 2--8 TB NVMe SSDs with RAID ²
NIC	10GbE	NVIDIA ConnectX®-6 Dx (25GbE)	NVIDIA ConnectX-6 Dx (25GbE)
OS	Ubuntu/RHEL/SUSE ¹	Ubuntu/RHEL/SUSE ¹	Ubuntu/RHEL/SUSE ¹
GPU	RTX PRO 6000 Blackwell Workstation Edition or RTX PRO 6000 Blackwell Max-Q Workstation Edition	2x RTX PRO 6000 Blackwell Max-Q Workstation Edition	4x RTX PRO 6000 Blackwell Max-Q Workstation Edition



Laptop	Best
CPU	Intel Core Ultra 9 Series
System Memory	64 GB DDR5
Storage	1 TB
OS	Ubuntu/RHEL/SUSE ¹
GPU	NVIDIA RTX PRO 5000 Blackwell Generation

NVIDIA RTX PRO AI Workstations

1. See NVIDIA Enterprise documentation for exact release support

2. For large datasets that need fast/reliable storage

Enterprise-Ready Solutions

Available from leading OEM manufacturers, NVIDIA RTX PRO AI workstations are designed and built for demanding enterprise deployments. Powered by the latest generation of workstation CPUs and available with NVIDIA ConnectX high-performance networking solutions, NVIDIA AI workstations are ready to tackle demanding AI development workflows. The latest generation of OEM desktop, mobile, and virtual workstations are available now.

With a full stack of enterprise-level deployment, support, and optimization tools, NVIDIA AI workstations easily fit into existing IT infrastructure, providing drop-in solutions for AI training, development, and inferencing on the desktop or in the data center. The NVIDIA GPU architecture scales from cloud to data center, desktop, laptop, and embedded devices, supporting the same software stack across devices, which enables users to move workloads seamlessly between them.

Ready to Get Started?

To learn more about NVIDIA RTX PRO AI Workstations, visit:
nvidia.com/ai-workstations

Contact sales at: <https://hypertecsp.com/contact/>

© 2025 NVIDIA Corporation and affiliates. All rights reserved. NVIDIA, the NVIDIA logo, ConnectX, CUDA, CUDA-X, RTX, and RTX PRO are trademarks and/or registered trademarks of NVIDIA Corporation and affiliates in the U.S. and other countries. Other company and product names may be trademarks of the respective owners with which they are associated. 4565200. DEC25

