

RAPIDS

GPU-Accelerated Data Analytics & Machine Learning

[Home](#) > [Deep Learning](#) > RAPIDS

The RAPIDS suite of software libraries, built on CUDA-X AI, gives you the freedom to execute end-to-end data science and analytics pipelines entirely on GPUs. It relies on NVIDIA® CUDA® primitives for low-level compute optimization, but exposes that GPU parallelism and high-bandwidth memory speed through user-friendly Python interfaces.

RAPIDS also focuses on common data preparation tasks for analytics and data science. This includes a familiar DataFrame API that integrates with a variety of machine learning algorithms for end-to-end pipeline accelerations without paying typical serialization costs. RAPIDS also includes support for multi-node, multi-GPU deployments, enabling vastly accelerated processing and training on much larger dataset sizes.



[RAPIDS Webpage](#)

[Intro Blog](#)

Features

<p>Hassle-Free Integration Accelerate your Python data science toolchain with minimal code changes and no new tools to learn.</p>	<p>Top Model Accuracy Increase machine learning model accuracy by iterating on models faster and deploying them more frequently.</p>
<p>Reduced Training Time Drastically improve your productivity with near-interactive data science.</p>	<p>Open Source Customizable, extensible, interoperable - the open-source software is supported by NVIDIA and built on Apache Arrow.</p>

Get Started

RAPIDS libraries are open source, written in Python, and built on Apache Arrow. The software is being developed in partnership with enterprises globally. Download RAPIDS to dramatically accelerate machine learning and data science.

[RAPIDS Webpage](#)