

Accelerating End-to-End Data Science Workflows (NVIDIA)

Duration: 08 hours

In this course, developers will learn how to build and execute end-to-end GPU accelerated data science workflows that enable them to quickly explore, iterate, and get their work into production.

About this Course

Whether you work at a software company that needs to improve customer retention, a financial services company that needs to mitigate risk, or a retail company interested in predicting customer purchasing behavior, your organization is tasked with preparing, managing, and gleaning insights from large volumes of data without wasting critical resources. Traditional CPU-driven data science workflows can be cumbersome, but with the power of GPUs, your teams can make sense of data quickly to drive business decisions.

Learning Objectives

In this course, developers will learn how to build and execute end-to-end GPU accelerated data science workflows that enable them to quickly explore, iterate, and get their work into production. Using the RAPIDS accelerated data science libraries, developers will apply a wide variety of GPU-accelerated machine learning algorithms, including XGBoost, cuGRAPH's single-source shortest path, and cuML's KNN, DBSCAN, and logistic regression to perform data analysis at scale.

By participating in this course, you'll learn how to:

- 📄 Implement GPU-accelerated data preparation and feature extraction using cuDF and Apache Arrow data frames.
- 📄 Apply a broad spectrum of GPU-accelerated machine learning tasks using XGBoost and a variety of cuML algorithms.
- 📄 Execute GPU-accelerated graph analysis with cuGraph, achieving massive-scale analytics in small amounts of time.
- 📄 Build beautiful data visualizations with the GPU-accelerated cuXFilter.

Upon completion, participants will be able to load, manipulate, and analyze data orders of magnitude fast, enabling more iteration cycles and drastically improving productivity.

Topics Covered

In this self-paced course, developers will learn how to build and execute end-to-end GPU accelerated data science workflows that enable them to quickly explore, iterate, and get their work into production. Using the RAPIDS accelerated data science libraries, developers will apply a wide variety of GPU-accelerated machine learning algorithms, including XGBoost, cuGRAPH's single-source shortest path, and cuML's KNN, DBSCAN, and logistic regression to perform data analysis at scale.

Course Outline

- Implement GPU-accelerated data preparation and feature extraction using cuDF and Apache Arrow data frames.
- Apply a broad spectrum of GPU-accelerated machine learning tasks using XGBoost and a variety of cuML algorithms.
- Execute GPU-accelerated graph analysis with cuGraph, achieving massive-scale analytics in small amounts of time.
- Build beautiful data visualizations with the GPU-accelerated cuXFilter.