

Build machine learning solutions using Azure Databricks

Courses Code: DP-3014

Duration: 1 Day

Delivery Mode: Instructor-led Training (ILT) | Virtual Instructor led Training (VILT)

Overview

Azure Databricks is a fully managed, cloud-based data analytics platform, which empowers developers to accelerate AI and innovation by simplifying the process of building enterprise-grade data applications. Built as a joint effort by Microsoft and the team that started Apache Spark, Azure Databricks provides data science, engineering, and analytical teams with a single platform for big data processing and machine learning. In this course, you'll learn how to use Azure Databricks to train and deploy machine learning models.

Skills Covered

- Use Azure Databricks to train and deploy machine learning models.
- Explore Azure Databricks and identify workloads, concepts, and governance features (such as Unity Catalog, Microsoft Purview).
- Use Apache Spark within Azure Databricks for data processing, analytics, and visualization.
- Train a machine learning model in Azure Databricks: data preparation, training, evaluation.
- Use MLflow within Azure Databricks to log experiments, register and serve models.
- Tune hyperparameters at scale using Azure Databricks (e.g., Hyperopt).
- Use AutoML in Azure Databricks (via UI or code) for streamlined model development.
- Train deep learning models using Azure Databricks (e.g., PyTorch, distributed training).
- Manage machine learning solutions in production: automate pipelines, handle model versioning, deploy and manage models at scale in Azure Databricks.

Who Should Attend

This course is designed for aspiring data scientists and AI engineers who need to train and manage machine learning models by using Azure Databricks.

Prerequisites

This learning path assumes that you have experience of using Python to explore data and train machine learning models with common open source frameworks, like Scikit-Learn, PyTorch, and TensorFlow.

Course Modules

1.Explore Azure Databricks

Azure Databricks is a cloud service that provides a scalable platform for data analytics using Apache Spark.



2. Use Apache Spark in Azure Databricks

Azure Databricks is built on Apache Spark and enables data engineers and analysts to run Spark jobs to transform, analyze and visualize data at scale.

3. Train a machine learning model in Azure Databricks

Machine learning involves using data to train a predictive model. Azure Databricks support multiple commonly used machine learning frameworks that you can use to train models.

4. Use MLflow in Azure Databricks

MLflow is an open source platform for managing the machine learning lifecycle that is natively supported in Azure Databricks.

5. Tune hyperparameters in Azure Databricks

Tuning hyperparameters is an essential part of machine learning. In Azure Databricks, you can use the Hyperopt library to optimize hyperparameters automatically.

6. Use AutoML in Azure Databricks

AutoML in Azure Databricks simplifies the process of building an effective machine learning model for your data.

7. Train deep learning models in Azure Databricks

Deep learning uses neural networks to train highly effective machine learning models for complex forecasting, computer vision, natural language processing, and other AI workloads.

8. Manage machine learning in production with Azure Databricks

Machine learning enables data-driven decision-making and automation, but deploying models into production for real-time insights is challenging. Azure Databricks simplifies this process by providing a unified platform for building, training, and deploying machine learning models at scale, fostering collaboration between data scientists and engineers.

Ref: <https://trainocate.com/courses/microsoft/dp-3014-implement-a-machine-learning-solution-with-azure-databricks-training>

END OF THE PAGE

