

Day - 1

• Overview of Databricks Platform

- Introduction to Databricks: Architecture, features, and key components
- Databricks Workspaces: Notebooks, Repos, and Jobs
- Databricks Runtime: Spark, MLflow, and integration with Apache Spark
- Setting up your environment and clusters for machine learning
- Hands-on Lab: Creating a Databricks workspace and setting up clusters.

• Introduction to Machine Learning with Databricks

- Introduction to Machine Learning concepts: Supervised and unsupervised learning
- Overview of MLflow: Model tracking, experimentation, and versioning
- Data preparation and feature engineering in Databricks
- Hands-on Lab: Loading datasets, exploring and preparing data for ML.

• Supervised Learning Algorithms

- Regression vs Classification: Key differences and use cases
- Implementing Linear Regression and Decision Trees using Spark MLlib
- Hands-on Lab: Building a simple regression model using PySpark and Databricks.

• Evaluating Model Performance

- Evaluation Metrics: Accuracy, Precision, Recall, F1-Score, AUC
- Cross-validation and hyperparameter tuning
- Hands-on Lab: Evaluating and tuning a classification model (Random Forest or Logistic Regression)

- **Unsupervised Learning and Clustering**

- Introduction to Unsupervised Learning: Clustering, Dimensionality Reduction
- K-Means Clustering and Principal Component Analysis (PCA)
- Hands-on Lab: Implementing K-Means clustering for customer segmentation.

- **Advanced Machine Learning Techniques**

- Gradient Boosting Machines (GBMs) and Random Forests
- Introduction to XGBoost and LightGBM
- Feature engineering for better model performance
- Hands-on Lab: Building and tuning an XGBoost model in Databricks.

- **Introduction to MLflow for Model Tracking**

- Overview of MLflow: Tracking experiments and managing models
- Tracking parameters, metrics, and models in Databricks
- Saving, loading, and versioning models in MLflow
- Hands-on Lab: Tracking an experiment with MLflow and storing model artifacts.

- **Model Deployment Using Databricks**

- Deploying ML models as REST APIs on Databricks
- Introduction to Databricks Model Serving and serving models at scale
- Hands-on Lab: Deploying a trained ML model using MLflow and serving it as an API.

- **Building End-to-End Machine Learning Pipelines**

- Overview of an ML pipeline: Data processing, feature engineering, model building, evaluation
- Using Databricks Workflows for pipeline orchestration
- Automating and scheduling ML jobs with Databricks
- Hands-on Lab: Building an end-to-end machine learning pipeline in Databricks using Spark and MLlib.

- **Hyperparameter Tuning and Model Optimization**

- Introduction to hyperparameter optimization with GridSearchCV and RandomSearchCV
- Automated machine learning with Databricks AutoML
- Hands-on Lab: Tuning the hyperparameters of a classification model (e.g., Logistic Regression or Random Forest).

- **Model Interpretability and Feature Importance**

- Techniques for understanding model predictions: Feature importance, SHAP, and LIME
- Introduction to interpretability libraries in Python (SHAP, LIME)
- Hands-on Lab: Interpreting a model's decision with SHAP in Databricks.

- **Databricks Certified Machine Learning Associate Exam Preparation**

- Exam structure and preparation tips
- Review of the exam topics and sample questions
- Mock Exam: 30–40 questions simulating the actual certification exam.