



Practical Data Science with Amazon SageMaker

Practical Data Science with Amazon SageMaker

You will learn how to solve a real-world use case with Machine Learning (ML) and produce actionable results using Amazon SageMaker. This course walks through the stages of a typical data science process for Machine Learning from analyzing and visualizing a dataset to preparing the data, and feature engineering. Individuals will also learn practical aspects of model building, training, tuning, and deployment with Amazon SageMaker. Real life use case includes customer retention analysis to inform customer loyalty programs.

Course Contents

- Module 1: Introduction to machine learning
 - Module 2: Introduction to data prep and SageMaker
 - Module 3: Problem formulation and dataset
 - Module 4: Data analysis and visualization
 - Module 5: Training and evaluating a model
 - Module 6: Automatically tune a model
 - Module 7: Deployment / production readiness
 - Module 8: Relative cost of errors
 - Module 9: Amazon SageMaker architecture and features
- Accessing Amazon SageMaker notebooks in a VPC

You have access to the labs for another 14 days after the course. This way you can repeat exercises or deepen them individually.

E-Book You will receive the original course documentation by Amazon Web Services as an e-book.

Target Group

This course is intended for:

- Developers
- Data Scientists

Prerequisites

We recommend that attendees of this course have:
Familiarity with Python programming language
Basic understanding of Machine Learning

This Course in the Web



You can find the up-to-date information and options for ordering under the following link:
www.expertech-training.com/go/AWPD

Reservation

On our Website, you can reserve a course seat for 7 days free of charge and in a non-committal manner. This can also be done by phone under +49 6074/4868-0.

Guaranteed Course Dates

To ensure reliable planning, we are continuously offering a wide range of guaranteed course dates.

Your Tailor-Made Course!

We can precisely customize this course to your project and the corresponding requirements.

Training	Prices, excl. of V.A.T.	
Classes in Germany	1 Day	€ 775
Online Training	1 Day	€ 775
Date/course venue	Course language German	
10/06-10/06/24	10/06-10/06/24	

Status 03/08/2024



Table of Contents

Practical Data Science with Amazon SageMaker

Module 1: Introduction to machine learning	Amazon SageMaker Ground Truth
Types of ML	Amazon SageMaker Neo
Job Roles in ML	
Steps in the ML pipeline	
Module 2: Introduction to data prep and SageMaker	
Training and test dataset defined	
Introduction to SageMaker	
Demonstration: SageMaker console	
Demonstration: Launching a Jupyter notebook	
Module 3: Problem formulation and dataset preparation	
Business challenge: Customer churn	
Review customer churn dataset	
Module 4: Data analysis and visualization	
Demonstration: Loading and visualizing your dataset	
Exercise 1: Relating features to target variables	
Exercise 2: Relationships between attributes	
Demonstration: Cleaning the data	
Module 5: Training and evaluating a model	
Types of algorithms	
XGBoost and SageMaker	
Demonstration: Training the data	
Exercise 3: Finishing the estimator definition	
Exercise 4: Setting hyper parameters	
Exercise 5: Deploying the model	
Demonstration: hyper parameter tuning with SageMaker	
Demonstration: Evaluating model performance	
Module 6: Automatically tune a model	
Automatic hyper parameter tuning with SageMaker	
Exercises 6-9: Tuning jobs	
Module 7: Deployment / production readiness	
Deploying a model to an endpoint	
A/B deployment for testing	
Auto Scaling	
Demonstration: Configure and test auto scaling	
Demonstration: Check hyper parameter tuning job	
Demonstration: AWS Auto Scaling	
Exercise 10-11: Set up AWS Auto Scaling	
Module 8: Relative cost of errors	
Cost of various error types	
Demo: Binary classification cutoff	
Module 9: Amazon SageMaker architecture and features	
Accessing Amazon SageMaker notebooks in a VPC	
Amazon SageMaker batch transforms	

