Event description

Are you interested in machine learning, but not sure where to start? Join us for this session with an AWS expert and demystify the basics. Using real-world examples, you'll learn about important concepts, terminology, and the phases of a machine learning pipeline. Learn how you can unlock new insights and value for your business using machine learning.

- Level: Fundamental
- Duration: 1.5 hours

Key topics covered

During this event, you will learn:

- What is Machine Learning?
- What is the machine learning pipeline, and what are its phases?
- What is the difference between supervised and unsupervised learning?
- What is reinforcement learning?
- What is deep learning?

Intended audience

This event is intended for:

- Developers
- Solution architects
- Data engineers
- Individuals interested in building solutions with machine learning - no machine learning experience required!

Recommended follow-up training and resources

We recommend that attendees of this event continue learning with these:

- Courses
  - o Deep Learning on AWS
  - o MLOps Engineering on AWS
  - o Practical Data Science with Amazon SageMaker
  - o The Machine Learning Pipeline on AWS

Event Outline

Section 1: Machine learning basics

- Classical programming vs. machine learning approach
- What is a model?
- Algorithm features, weights, and outputs
- Machine learning algorithm categories
- Supervised algorithms

- Unsupervised algorithms
- Reinforcement learning

Section 2: What is deep learning?
- How does deep learning work?
- How deep learning is different

Section 3: The Machine Learning Pipeline
- Overview
- Business problem
- Data collection and integration
- Data processing and visualization
- Feature engineering
- Model training and tuning
- Model evaluation
- Model deployment

Section 4: What are my next steps?
- Resources to continue learning