

AI-300T00: Operationalize machine learning and generative AI solutions

Duration: 32 hours

Course Overview

This course prepares learners to design, implement, and operate Machine Learning Operations (MLOps) and Generative AI Operations (GenAI Ops) solutions on Azure. It covers building secure and scalable AI infrastructure, managing the full lifecycle of traditional machine learning models with Azure Machine Learning, and deploying, evaluating, monitoring, and optimizing generative AI applications and agents using Microsoft Foundry. Learners will gain hands-on knowledge of automation, continuous integration and delivery, infrastructure as code, and observability by using tools such as GitHub Actions, Azure CLI, and Bicep. The course emphasizes collaboration with data science and DevOps teams to deliver reliable, production-ready AI systems aligned with modern MLOps and GenAI Ops best practices.

Audience Profile

This course is intended for data scientists, machine learning engineers, and DevOps professionals who want to design and operate production-grade AI solutions on Azure. It is suited for learners with experience in Python, a foundational understanding of machine learning concepts, and basic familiarity with DevOps practices such as source control, CI/CD, and command-line tools, who are preparing to implement MLOps and GenAI Ops workflows using Azure-native services.

Session Agenda

Learning path 1: Operationalize machine learning models (MLOps)

Module 1: Experiment with Azure Machine Learning

- Introduction
- Preprocess data and configure featurization
- Run an automated machine learning experiment
- Evaluate and compare models
- Configure MLflow for model tracking in notebooks
- Train and track models in notebooks
- Evaluate models with the Responsible AI dashboard

Module 2: Perform hyperparameter tuning with Azure Machine Learning

- Introduction
- Define a search space
- Configure a sampling method

- Configure early termination
- Use a sweep job for hyperparameter tuning

Module 3: Run pipelines in Azure Machine Learning

- Create Components
- Create a Pipeline
- Run a Pipeline job

Module 4: Trigger Azure Machine Learning jobs with GitHub Actions

- Understand the business problem
- Explore the solution architecture
- Use Github Actions for model training

Module 5: Trigger GitHub Actions with feature-based development

- Understand the business problem
- Explore the solution architecture
- Trigger a workflow

Module 6: Work with environments in GitHub Actions

- Understand the business problem
- Explore the solution architecture
- Set up environments

Module 7: Deploy a model with GitHub Actions

- Understand the business problem
- Explore the solution architecture
- Model Deployment

Learning path 2: Operationalize generative AI applications (GenAIOps)

Module 1: Plan and prepare a GenAIOps solution

- Explore use cases for GenAIOps
- Select the right generative AI model
- Understand the development lifecycle of a language model application
- Explore available tools and frameworks to implement GenAIOps

Module 2: Manage prompts for agents in Microsoft Foundry with GitHub

- Apply version control to prompts
- Understand Microsoft Foundry agents and prompt versioning
- Organize prompts in GitHub repositories
- Develop safe prompt deployment workflows

Module 3: Evaluate and optimize AI agents through structured experiments

- Design evaluation experiments
- Apply Git-based workflows to optimization experiments
- Apply evaluation rubrics for consistent scoring

Module 4: Automate AI evaluations with Microsoft Foundry and GitHub Actions

- Understand why automated evaluations matter
- Align evaluators with human criteria
- Create evaluation datasets
- Implement batch evaluations with Python
- Integrate evaluations into GitHub Actions

Module 5: Monitor your generative AI application

- Why do you need to monitor?
- Understand key metrics to monitor
- Explore how to monitor with Azure
- Integrate monitoring into your app
- Interpret monitoring results

Module 6: Analyze and debug your generative AI app with tracing

- Why do you need to use tracing?
- Identify what to trace in generative AI applications
- Implement tracing in generative AI applications
- Debug complex workflows with advanced tracing patterns
- Make informed decisions with trace data analysis