Hewlett Packard Enterprise

Building Artificial Intelligence Models Workshop

H38HQS

Course ID	H38HQS
Duration	2 days
Format	ILT, VILT
View schedule, local pricing, and register	View now
Browse related courses	View now

This course presents an enterprise approach to understand, evaluate, and apply MLOps and artificial intelligence (AI) solutions to a business problem. The course first focuses on forming AI components—notably computer vision and natural language processing—and their applications. You then build a practical architecture that can be deployed for public consumption using Docker containers. This course also helps you decide when to select pre-built models, use custom models, or design your own.

5 reasons to choose HPE as your training partner

1. Learn HPE and in-demand IT industry technologies from expert instructors.
2. Build career-advancing power skills.
3. Enjoy personalized learning journeys aligned to your company's needs.
4. Choose how you learn: in-person, virtually, or online—anytime, anywhere.
5. Sharpen your skills with access to real environments in virtual labs.

Explore our simplified purchase options, including HPE Education Learning Credits.

Audience

This course is ideal for software engineers, IT professionals, data engineers, database professionals, developers, testers, solution architects, AI and automation enthusiasts, statisticians, and other professionals looking to build MLOps and AI capabilities.

Prerequisites

A basic understanding of any programming or scripting language is recommended.

Course objectives

After completing this course, you should be able to:

- Understand the fundamentals and key concepts of AI
- Develop and implement AI algorithms and models
- Build and deploy AI models for consumption
- Build reusable containers for AI deployments
- Understand the ethical challenges of AI
- Understand generative AI

Detailed course outline

Module 1: The Ecosystem and Overview	• What is artificial intelligence? • History of AI • Types of AI systems	• Tools and technologies • Hardware dependencies
Module 2: Ethics in AI	• Principles of ethical AI • Law and compliance of AI	• Learning from mistakes
Module 3: Machine Learning (ML) Overview	• Narrow vs. general AI • How AI differs from ML • Pre-built vs custom models	• What is an AI algorithm? • Building your first ML model
Module 4: Process Automation	• Registering a dataset • Automated machine learning	• Building your first AI endpoint
Module 5: AI Applications—Computer Vision	• Computer vision vs. image processing • Types of pre-built models for computer vision	• Object detection
Module 6: AI Applications—Natural Language Processing (NLP)	• An NLP pipeline • Types of pre-built models for NLP	• Sentiment analysis • Sensitive data redaction
Module 7: Containerized Applications	• Understanding a container • Docker and Kubernetes Introduction	• Building an AI Docker container • Deploying an AI Docker container
Module 8: MLOps	• An enterprise architecture for chatbots • Data drift	• MLOps best practices • Monitoring the AI service
Module 9: Introduction to Generative AI	• What is generative AI • Generative models in AI • Large language models (LLM)	• Applications of generative AI • Responsible generative AI

Learn more at

hpe.com/ww/learnAI

Follow us:

Hewlett Packard Enterprise

H38HQS B.00, April 2024