Generate Smarter Generative AI Outputs on Google Cloud

Duration: 01 days (08 hours)

Courseware: Unofficial

Module 1: Introduction to AI and Machine Learning on Google Cloud

This module introduces the AI and machine learning (ML) offerings on Google Cloud that build both predictive and generative AI projects. It explores the technologies, products, and tools available throughout the data-to-AI life cycle, encompassing AI foundations, development, and solutions. It aims to help data scientists, AI developers, and ML engineers enhance their skills and knowledge through engaging learning experiences and practical hands-on exercises.

Module 2: Introduction to Image Generation

This module introduces diffusion models, a family of machine learning models that recently showed promise in the image generation space. Diffusion models draw inspiration from physics, specifically thermodynamics. Within the last few years, diffusion models became popular in both research and industry. Diffusion models underpin many state-of-the-art image generation models and tools on Google Cloud. This module introduces you to the theory behind diffusion models and how to train and deploy them on Vertex AI.

Module 3: Vector Search and Embeddings

This module introduces Vertex AI Vector Search and describes how it can be used to build a search application with large language model (LLM) APIs for embeddings. The course consists of conceptual lessons on vector search and text embeddings, practical demos on how to build vector search on Vertex AI, and a hands-on lab.

Module 4: Inspect Rich Documents with Gemini Multimodality and Multimodal RAG

Complete the intermediate **Inspect Rich Documents with Gemini Multimodality and Multimodal RAG** skill badge to demonstrate skills in the following: using multimodal prompts to extract information from text and visual data, generating a video description, and retrieving extra information beyond the video using multimodality with Gemini; building metadata of documents containing text and images, getting all relevant text chunks, and printing citations by using Multimodal Retrieval Augmented Generation (RAG) with Gemini.
