Google Generative AI for Developers

Course Duration: 03 days

No. of Hours: 08 hours/day

Note: Client needs to have their own google cloud free subscription via their own credit card: https://cloud.google.com/free

==========

Chapter 01: Introduction to Image Generation

This chapter 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 course introduces you to the theory behind diffusion models and how to train and deploy them on Vertex AI.

==========

Chapter 02: Attention Mechanisms

This chapter will introduce you to the attention mechanism, a powerful technique that allows neural networks to focus on specific parts of an input sequence. You will learn how attention works, and how it can be used to improve the performance of a variety of machine learning tasks, including machine translation, text summarization, and question answering.

==========

Chapter 03: Encoder-Decoder Architecture

This chapter gives you a synopsis of the encoder-decoder architecture, which is a powerful and prevalent machine learning architecture for sequence-to-sequence tasks such as machine translation, text summarization, and question answering. You learn about the main components of the encoder-decoder architecture and how to train and serve these models. In the corresponding lab walkthrough, you'll code in TensorFlow a simple implementation of the encoder-decoder architecture for poetry generation from the beginning.

==========

Chapter 04: Transformer Model and BERT model

This chapter introduces you to the Transformer architecture and the Bidirectional Encoder Representations from Transformers (BERT) model. You learn about the main components of the Transformer architecture, such as the self-attention mechanism, and how it is used to build the BERT model. You also learn about the different tasks that BERT can be used for, such as text classification, question answering, and natural language inference.

==========

Chapter 05: Create Image Captioning Model

This chapter teaches you how to create an image captioning model by using deep learning. You learn about the different components of an image captioning model, such as the encoder and decoder, and how to train and evaluate your model. By the end of this course, you will be able to create your own image captioning models and use them to generate captions for images

==========

Chapter 06: Introduction to Generative AI Studio

This chapter introduces Generative AI Studio, a product on Vertex AI, that helps you prototype and customize generative AI models so you can use their capabilities in your applications. In this course, you learn what Generative AI Studio is, its features and options, and how to use it by walking through demos of the product. In the end, you will have a quiz to test your knowledge.

==========

Chapter 07: Generative AI Explorer – Vertex AI

This chapter is a collection of labs on how to use Generative AI on Google Cloud. Through the labs, you will learn about how to use the models in the Vertex AI PaLM API family, including text-bison, chat-bison, and text embedding-gecko. You will also learn about prompt design, best practices, and how it can be used for ideation, text classification, text extraction, text summarization, and more. You will also learn how to tune a foundation model by training it via Vertex AI custom training and deploy it to a Vertex AI endpoint.

==========

Chapter 08: Explore and Evaluate Model using Model Garden

Model Garden on Vertex AI provides a single place to search, discover, and interact with a wide variety of models from Google and Google partners. Model Garden is available on Vertex AI and can be accessed from the Google Cloud console.

This chapter will help you explore lab for a variety of use cases to explore Model Garden and then use Generative AI Studio to create and experiment with prompts.

==========

Chapter 08: Prompt Design using PaLM

Prompt design is the process of creating prompts that are effective in generating the desired output from a large language model (LLM) like PaLM. Prompts can be used to generate text, translate languages, write different kinds of creative content, and answer your questions in an informative way.

To get a good result from your prompt, the user needs to be specific, use keywords that are relevant to the output you want to generate, and provide PaLM with examples of the output you want to help it to better understand what you are looking for.

In this chapter, you will have labs to create prompts that can extract summaries from very long home descriptions end users match their specific requirements.