Hewlett Packard Enterprise

NVIDIA Generative AI with Diffusion Models

H45WTS

Course ID	H45WTS
Duration	1 day
Format	ILT/VILT
Schedule, pricing & registration	View now
Browse related courses	View now

In this course, you will learn about the latest in generative artificial intelligence (AI). Improvements in computing power and scientific theory have made generative artificial intelligence more accessible than ever before. Generative AI plays a significant role across industries and will gain exponential importance due to the numerous applications it can be used for, such as creative content generation, data augmentation, simulation and planning, anomaly detection, drug discovery, and personalized recommendations. This course takes a deeper dive on denoising diffusion models, which are a popular choice for text-to-image pipelines, disrupting several industries.

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 on-demand—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 beginners with technical interest in generative AI and large language models (LLMs).

Prerequisites

Before attending this course, you should have a good understanding of PyTorch and deep learning

Course objectives

After completing this course, you should be able to:

- Build a U-Net to generate images from pure noise
- Improve the quality of generated images with the denoising diffusion process
- Compare Denoising Diffusion Probabilistic Models (DDPMs) with Denoising Diffusion Implicit Models (DDIMs)
- Control the image output with context embeddings
- Generate images from English text-prompts using CLIP

Detailed course outline

Module 1: Parameter-Efficient Fine-Tuning (PEFT) Essentials	• Build a U-Net, a type of autoencoder for images • Learn about transposed convolution to increase the size of an image	• Learn about non-sequential neural networks and residual connections • Experiment with feeding noise through the U-Net to generate new images
Module 2: Control with Context	• Learn how to alter the output of the diffusion process by adding context embeddings	• Add additional model optimizations such as Sinusoidal Position Embeddings, the GELU activation function, and Attention.
Module 3: Text-to-Image with CLIP	• Walk through the CLIP architecture to learn how it associates image embeddings with text embeddings	• Use CLIP to train a text-to-image diffusion model
Module 4: State-of-the-art Models	• Review various state-of-the-art generative AI models and connect them to the concepts learned in class • Discuss prompt engineering and how to better influence the output of generative AI models	• Learn about content authenticity and how to build trustworthy models

Learn more at

hpe.com/ww/learnAI

Follow us:

Hewlett Packard Enterprise

H45WTS A.00, October 2024