



Generative AI With Diffusion Models

NVIDIA Deep Learning Institute



Thanks to improvements in computing power and scientific theory, generative AI is more accessible than ever before. Generative AI will play a significant role across industries due to its numerous applications, such as creative content generation, data augmentation, simulation and planning, anomaly detection, drug discovery, personalized recommendations, and more. In this course, we'll take a deeper dive into denoising diffusion models, which are a popular choice for text-to-image pipelines, disrupting several industries.

Learning Objectives

In this workshop, you'll learn how 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 the Contrastive Language—Image Pretraining (CLIP) neural network

Overview

Duration	8 hours
Price	Contact us for pricing.
Prerequisites	> Good understanding of PyTorch. > Good understanding of deep learning.
Tools, libraries, and frameworks	PyTorch, CLIP.
Assessment type	The assessment will include: > Series of questions related to the lecture content. > A programming exercise in which students will be asked to create a denoising diffusion implicit model.
Certificate	Upon successful completion of the assessment, participants will receive an NVIDIA DLI certificate to recognize their subject matter competency and support professional career growth.
Hardware requirements	Desktop or laptop computer capable of running the latest version of Chrome or Firefox. Each participant will be provided with dedicated access to a fully configured, GPU-accelerated workstation in the cloud.
Language	English

Workshop Outline

Introduction

(15 minutes)

- > Meet the instructor.
- > Create an account at courses.nvidia.com/join.

From U-Nets to Diffusion

(60 minutes)

- > 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.

Break (10 minutes)

The Diffusion Process

(90 minutes)

- > Learn how to create a sequence of increasingly noisy images with the forward diffusion process (DDPM).
- > Learn how to skip ahead in the noise-generating sequence for faster network training (DDIM).
- > Use Bayes' theorem to remove noise from an image for the reverse diffusion process.

Break (60 minutes)

Control With Context

(60 minutes)

- > Learn how to alter the output of the diffusion process by adding context embeddings.
- > Add additional model optimizations such as:
 - > Sinusoidal position embeddings.
 - > The Gaussian Error Linear Unit (GELU) activation function.
 - > Attention.

Text-to-Image With CLIP

(60 minutes)

- > 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.

Break (10 minutes)

State-of-the-Art Models

(60 minutes)

- > 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.

Trustworthy AI

(60 minutes)

- > Learn more about content authenticity and how to build trustworthy models.
- > Review key learnings and answer questions.
- > Complete the assessment and earn a certificate.
- > Complete the workshop survey.

Next Steps

Continue learning with these DLI trainings:

- > [Building Conversational AI Applications](#)
- > [Data Parallelism: How to Train Deep Learning Models on Multiple GPUs](#)
- > [Building Transformer-Based Natural Language Processing Applications](#)

Why Choose NVIDIA Deep Learning Institute for Hands-On Training?

- > Access workshops from anywhere with just your desktop or laptop and an internet connection. Each participant will have access to a fully configured, GPU-accelerated server in the cloud.
- > Obtain hands-on experience with the most widely used, industry-standard software, tools, and frameworks.
- > Learn to build deep learning and accelerated computing applications for industries, such as healthcare, robotics, manufacturing, accelerated computing, and more.
- > Gain real-world expertise through content designed by NVIDIA and industry experts.
- > Earn an NVIDIA DLI certificate to demonstrate your subject matter competency and support your career growth.

Ready to Get Started?

For the latest DLI workshops and trainings, visit nvidia.com/dli

For questions, contact us at nvdl@nvidia.com