Introduction to GANs

Manu Mathew Thomas
Creative Coding Lab

Previously

Image Classifier

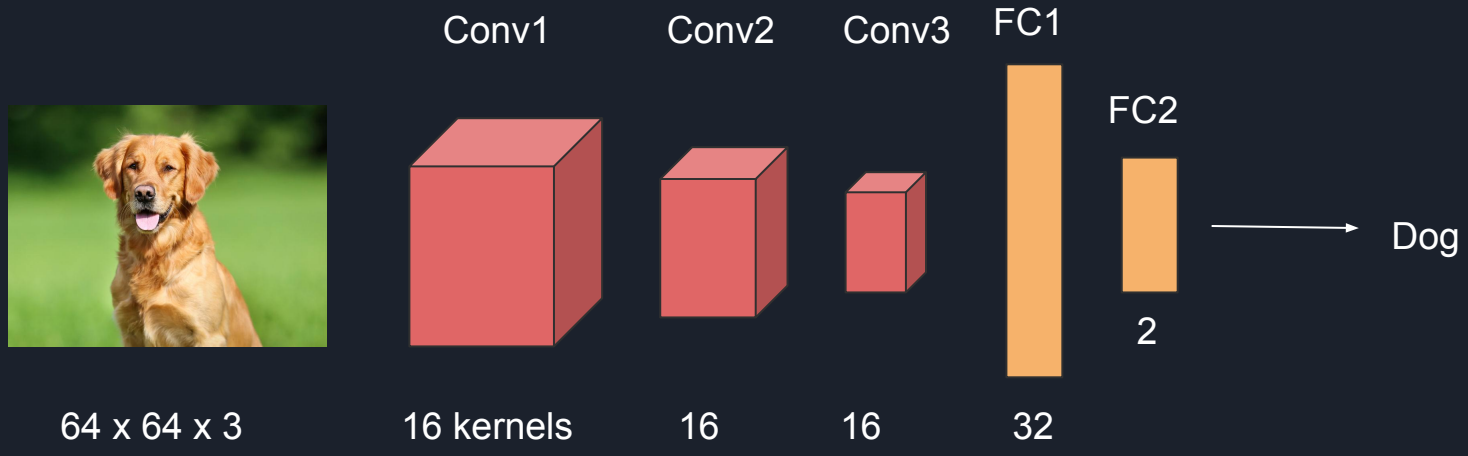

Previously

Image Generator

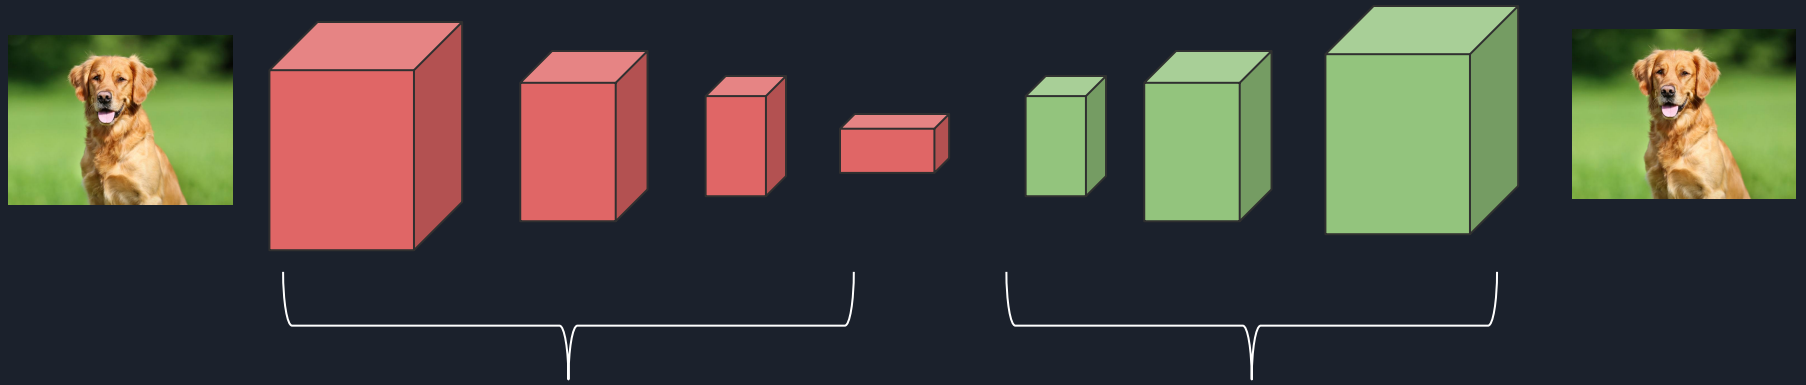

Generative Adversarial Networks

Image Generator

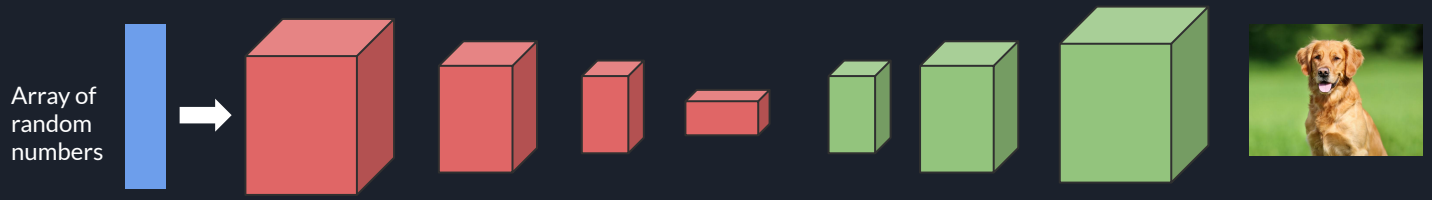

Image Classifier

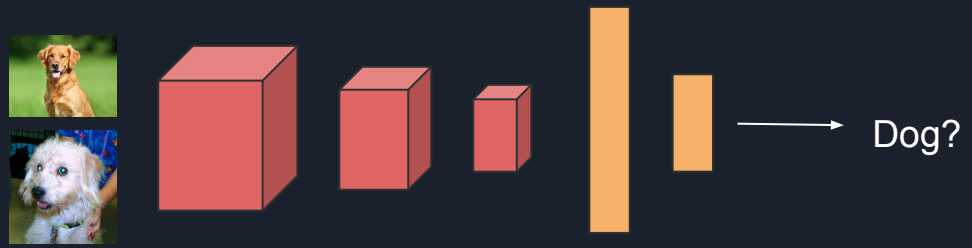

Generative Adversarial Networks

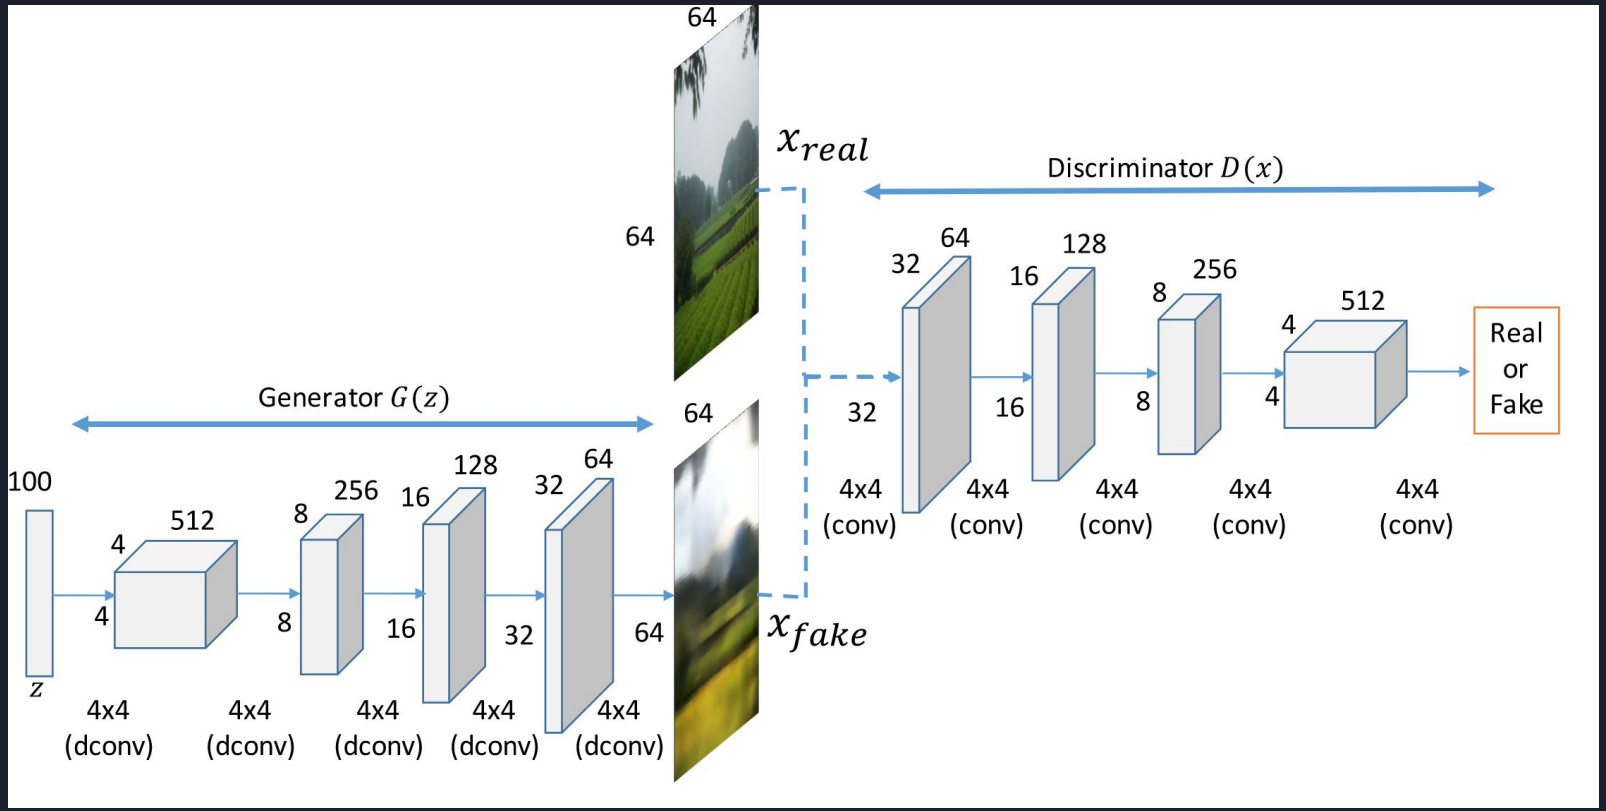

Conditional Generative Adversarial Networks

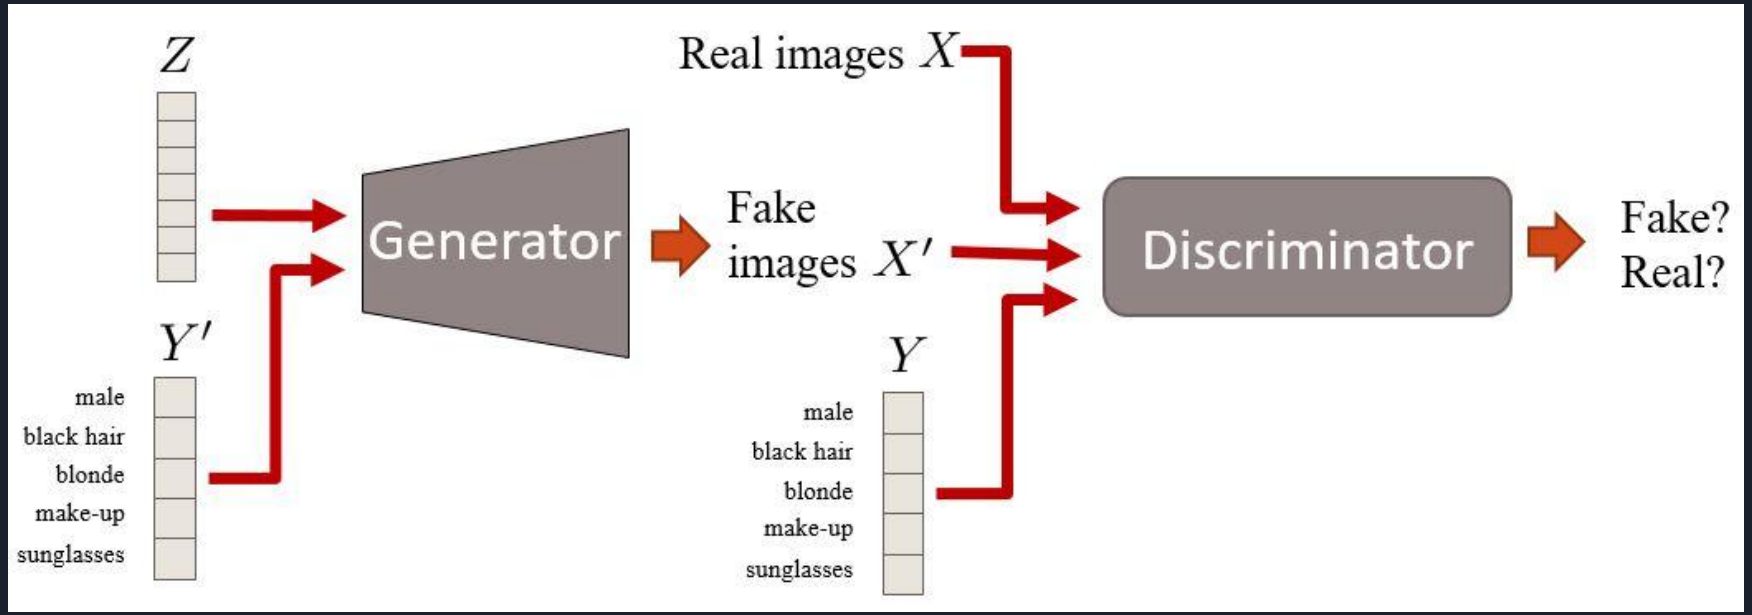

Conditional Generative Adversarial Networks

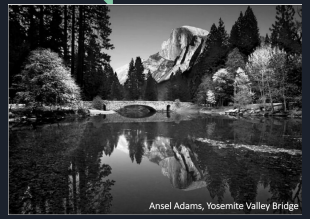

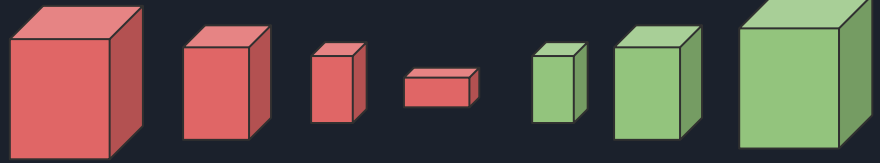

Generator

Real

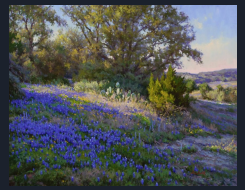

Or

Fake

True
or
False

Discriminator

GAN basics

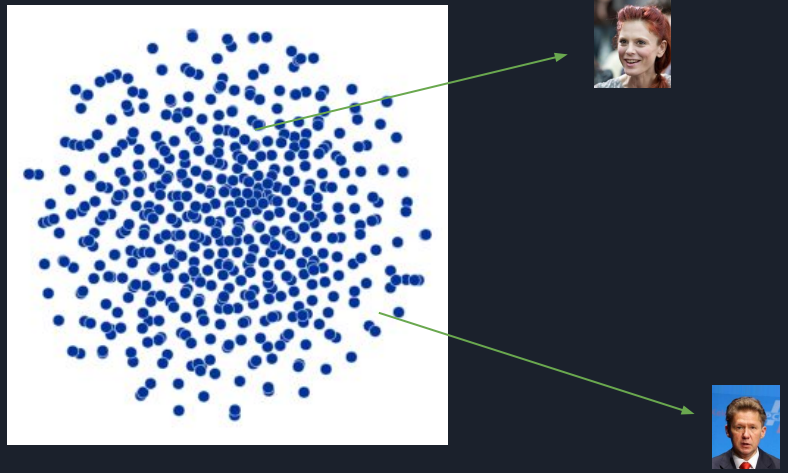

Distribution of human face images

~12K dimensions

Each image is 64x64 pixels = 4096 pixels

3 channels RGB; 4096 * 3 = 12,288 pixels

Each pixel is one dimension

A single point in this distribution corresponds to a vector of 12,228 pixels which can be reshaped to 64x64x3

This is a complex distribution and properties are unknown to us

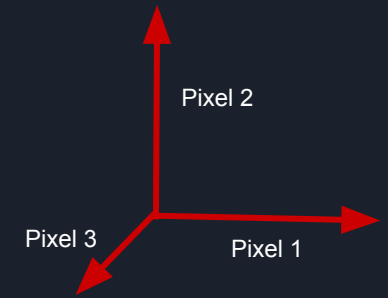

GAN basics

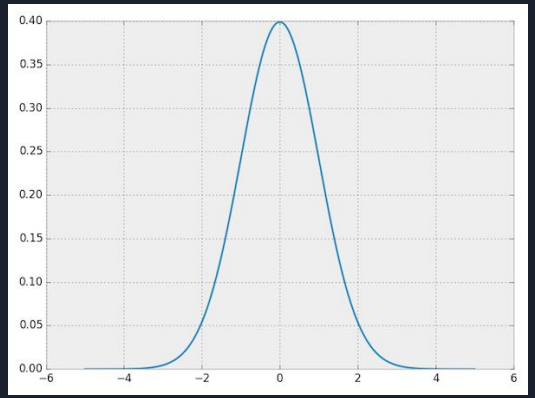

Gaussian Distribution

Simple and well-studied distribution

F(Gaussian noise) = point in the face image distribution space (P)

F(P) = true or false

F(Gaussian noise) is the generator
F(P) is the discriminator

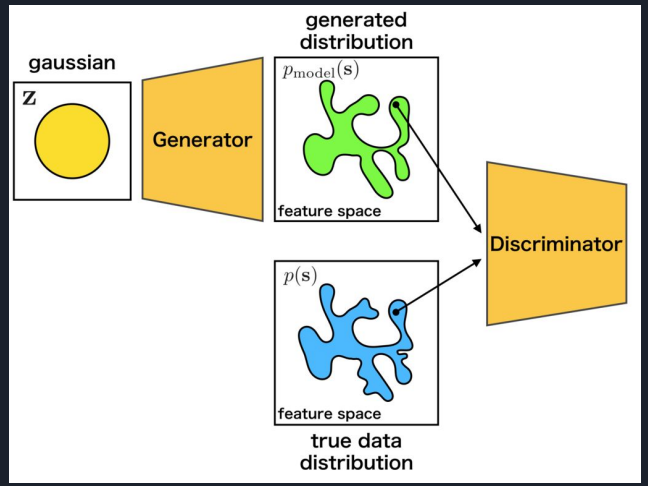

GAN Denoising

Input

Output

Ground Truth

GAN Denoising

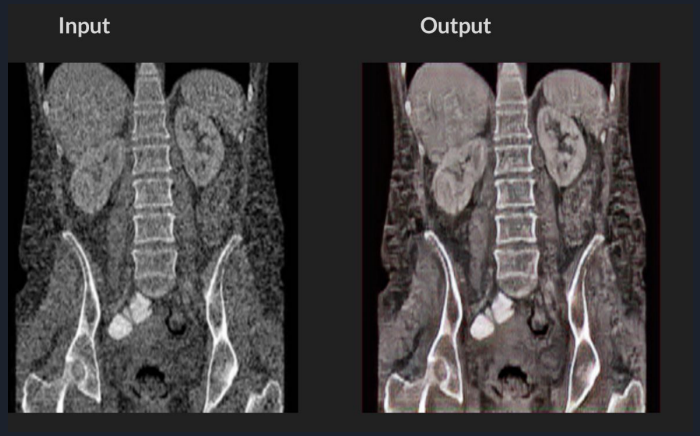

Input

Output

GAN Denoising

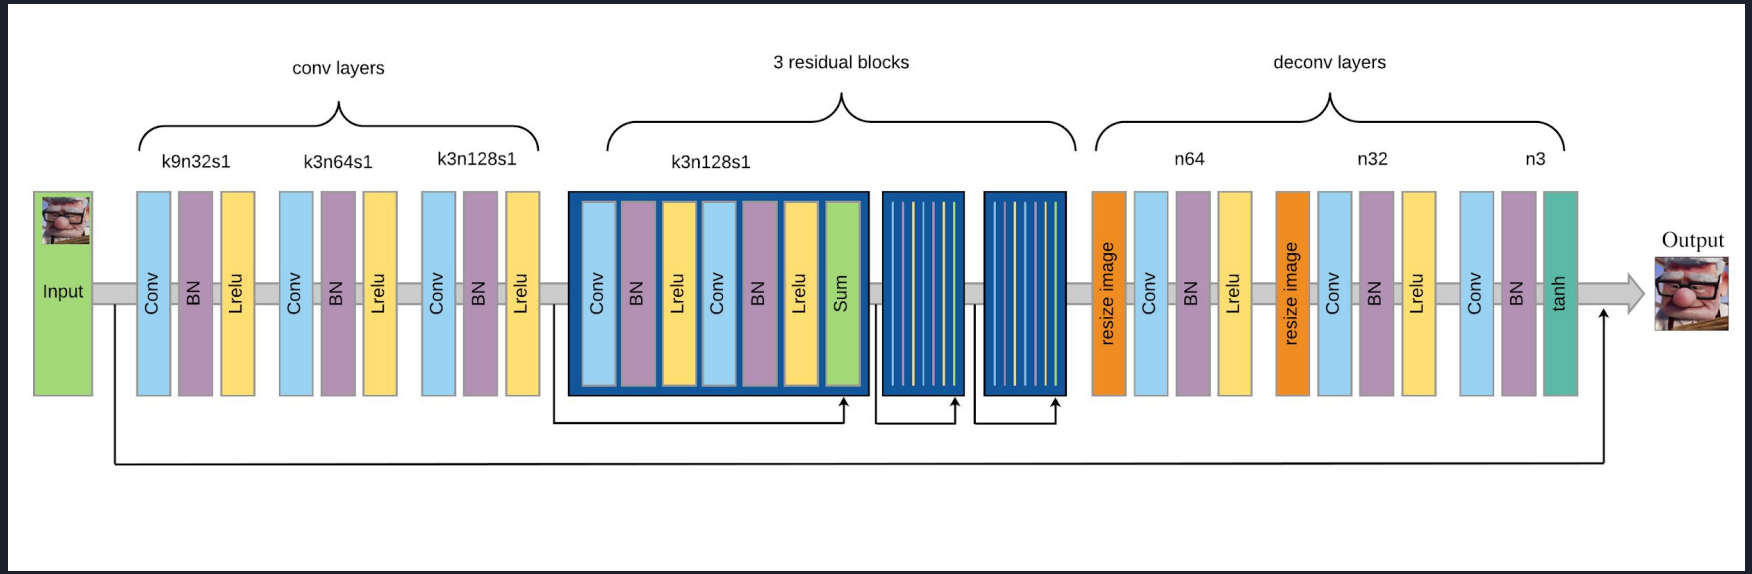

Generator

GAN Denoising

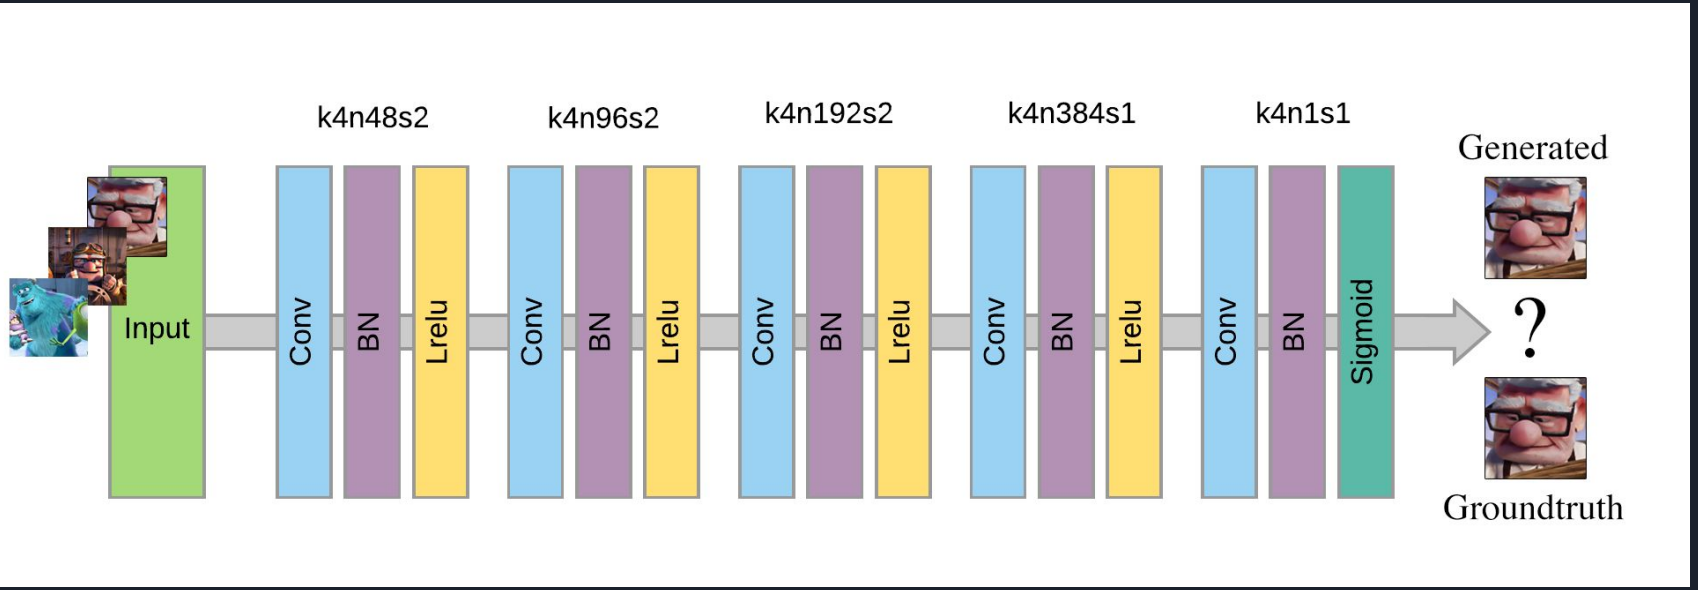

Deep Illumination - indirect illumination with GANs

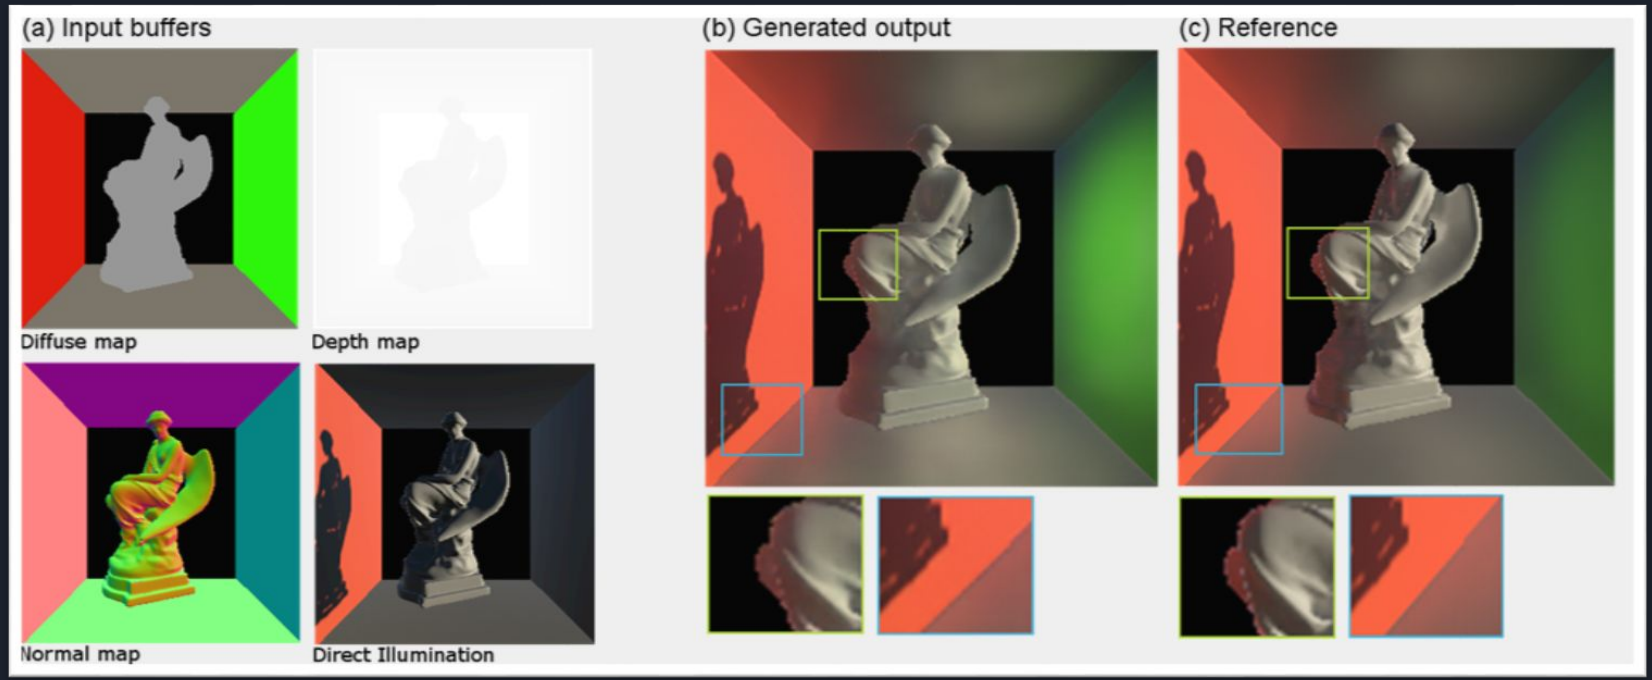

Deep Illumination - indirect illumination with GANs

Training set

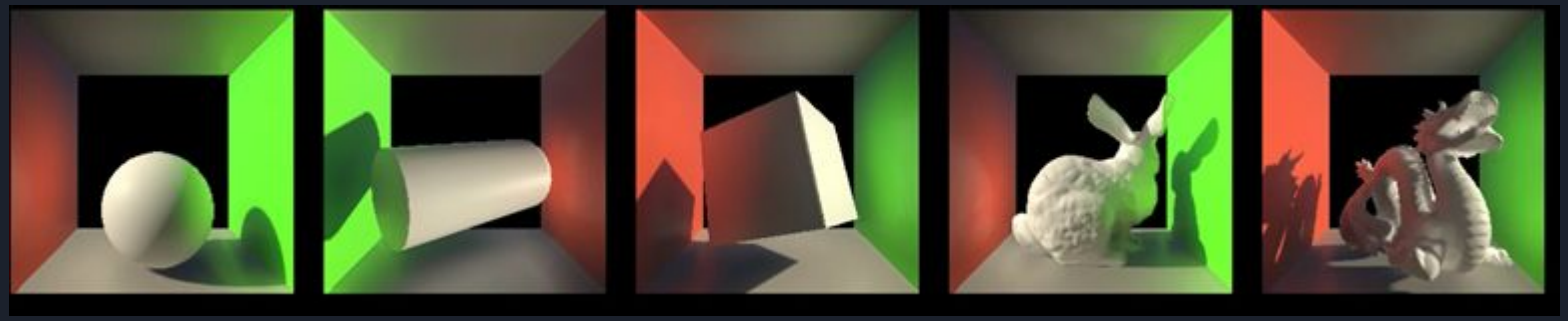

Test set

Deep Illumination - indirect illumination with GANs

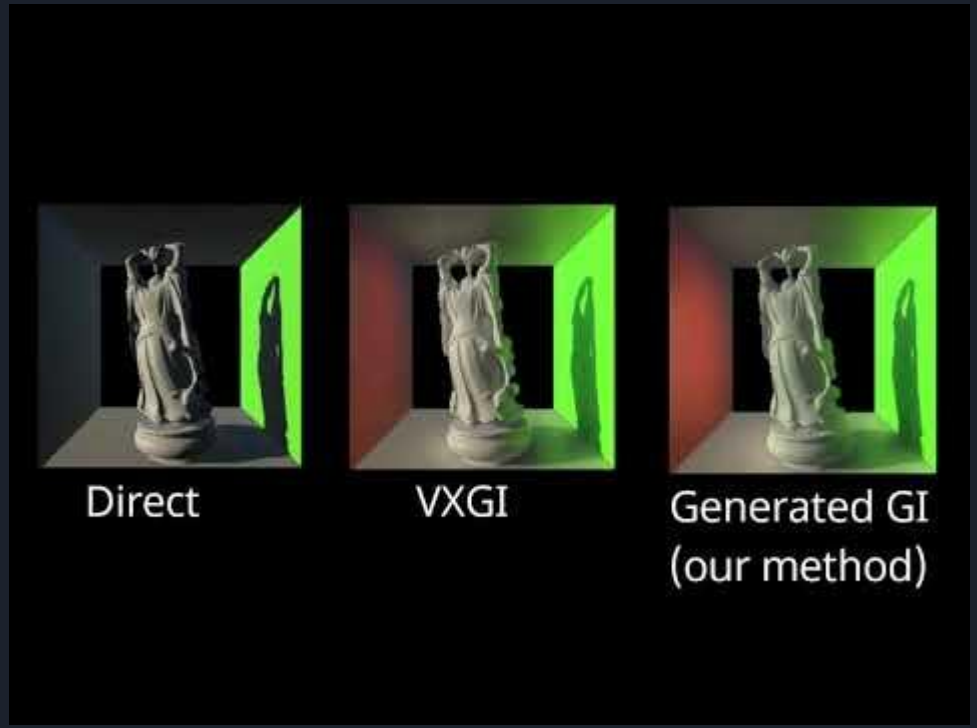

Deep Illumination - indirect illumination with GANs

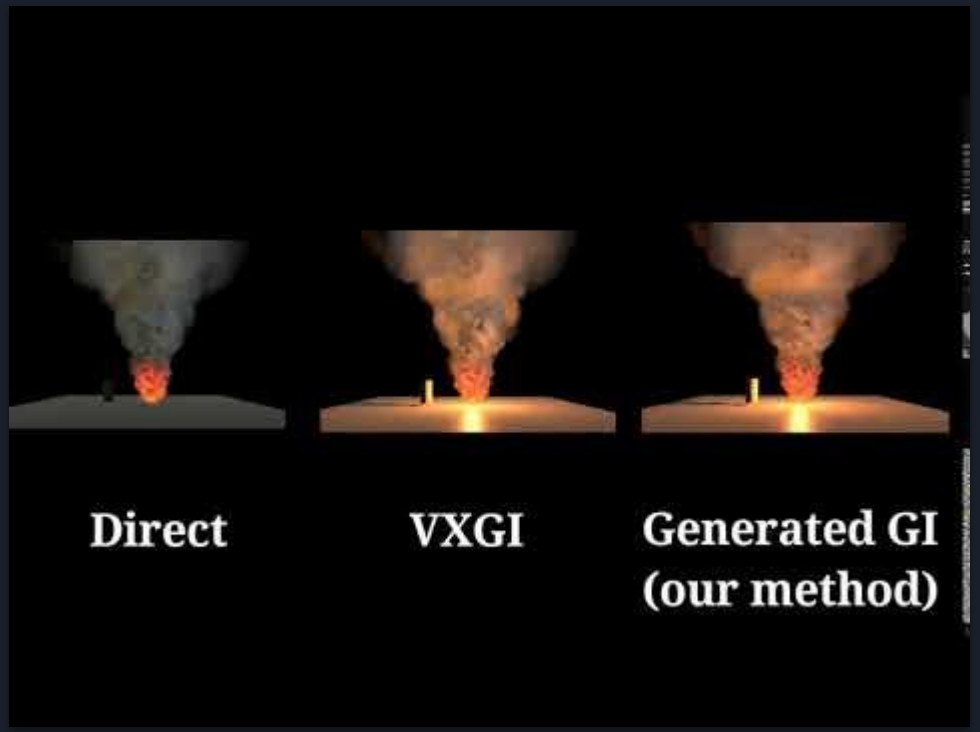

Deep Illumination - indirect illumination with GANs