

Deep Convolutional Generative Adversarial Networks

DEEP LEARNING BASED ANOMALY DETECTION MODELING (GEK6207.01-00)

Coding practice

Sunghyun Ahn

skd@yonsei.ac.kr

<2023/11/09>


1 DCGAN

What is a DCGAN?

A DCGAN is a direct extension of the GAN, except that it explicitly uses **convolutional and convolutional-transpose layers in the discriminator and generator**, respectively.


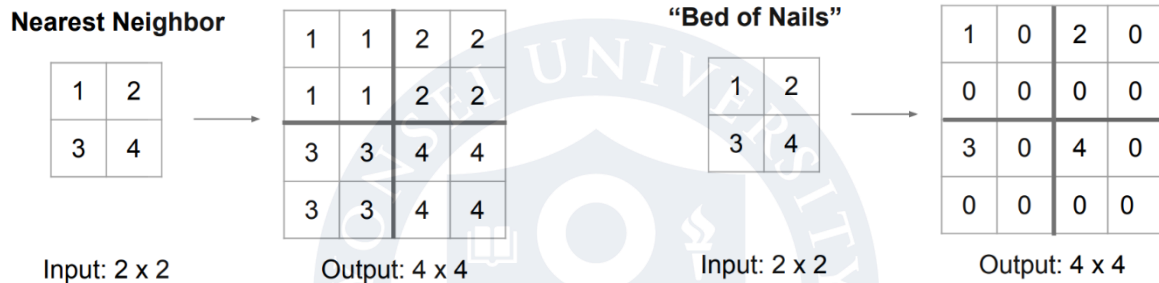
The **generator** is comprised of **convolutional-transpose layers, batch norm layers, and ReLU activations**. The input is a latent vector, that is drawn from a standard normal distribution and the output is a $3 \times 64 \times 64$ RGB image.

The **discriminator** is made up of **strided convolution layers, batch norm layers, and LeakyReLU activations**. The input is a $3 \times 64 \times 64$ input image and the output is a scalar probability that the input is from the real data distribution.




1 DCGAN

Upsampling




Corresponding pairs of downsampling and upsampling layers




1 DCGAN

Upsampling



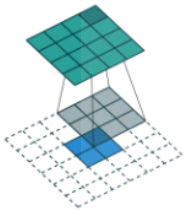
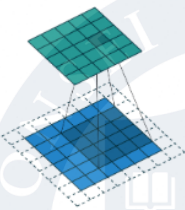
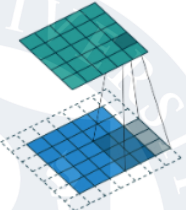
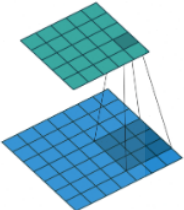
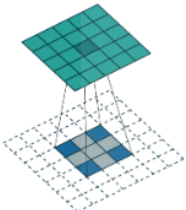
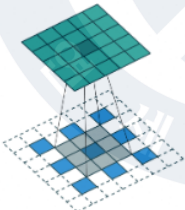
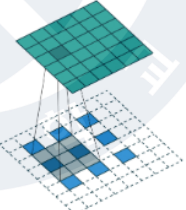
Linear Interpolation



Transposed Convolution

1 DCGAN


Transposed convolution

			
No padding, no strides, transposed	Arbitrary padding, no strides, transposed	Half padding, no strides, transposed	Full padding, no strides, transposed
			
No padding, strides, transposed	Padding, strides, transposed	Padding, strides, transposed (odd)	

https://github.com/vdumoulin/conv_arithmetic


1 DCGAN

Convolution output size



in = 5
out = 3
padding (p) = 1
stride (s) = 2
kernel_size (k) = 3

$$\text{Out} = \left\lfloor \frac{\text{in} + 2 * \text{padding} - \text{kernel_size}}{\text{stride}} \right\rfloor + 1$$




in = 3
out = 5
padding (p) = 1
stride (s) = 2
kernel_size (k) = 3

$$\text{Out} = (\text{in} - 1) * \text{stride} - 2 * \text{padding} + \text{kernel_size}$$

1 DCGAN

Celeb-A Faces dataset

Training Images



- Celeb-A Faces dataset is a large collection of celebrity facial images used for tasks like facial recognition and generative modeling.
- It includes over **200,000** diverse celebrity images.

1 DCGAN



Inputs Code

In [44]:

```
# Root directory for dataset
dataroot = "data/celeba"

# Number of workers for dataloader
workers = 2

# Batch size during training
batch_size = 128

# Spatial size of training images. All images will be resized to this
# size using a transformer.
image_size = 64

# Number of channels in the training images. For color images this is 3
nc = 3

# Size of z latent vector (i.e. size of generator input)
nz = 100

# Size of feature maps in generator
ngf = 64


# Size of feature maps in discriminator
ndf = 64

# Number of training epochs
num_epochs = 5

# Learning rate for optimizers
lr = 0.0002


# Beta hyperparameter for Adam optimizers
beta1 = 0.5

# Number of GPUs available. Use 0 for CPU mode.
ngpu = 1
```



1 DCGAN

Generator Code



In [47]:

```
# Generator Code

class Generator(nn.Module):
    def __init__(self, ngpu):
        super(Generator, self).__init__()
        self.ngpu = ngpu
        self.main = nn.Sequential(
            # input is Z, going into a convolution
            nn.ConvTranspose2d( nz, ngf * 8, 4, 1, 0, bias=False),
            nn.BatchNorm2d(ngf * 8),
            nn.ReLU(True),
            # state size. ``(ngf*8) x 4 x 4``
            nn.ConvTranspose2d(ngf * 8, ngf * 4, 4, 2, 1, bias=False),
            nn.BatchNorm2d(ngf * 4),
            nn.ReLU(True),
            # state size. ``(ngf*4) x 8 x 8``
            nn.ConvTranspose2d( ngf * 4, ngf * 2, 4, 2, 1, bias=False),
            nn.BatchNorm2d(ngf * 2),
            nn.ReLU(True),
            # state size. ``(ngf*2) x 16 x 16``
            nn.ConvTranspose2d( ngf * 2, ngf, 4, 2, 1, bias=False),
            nn.BatchNorm2d(ngf),
            nn.ReLU(True),
            # state size. ``(ngf) x 32 x 32``
            nn.ConvTranspose2d( ngf, nc, 4, 2, 1, bias=False),
            nn.Tanh()
            # state size. ``(nc) x 64 x 64``
        )

    def forward(self, input):
        return self.main(input)
```

```
# Number of channels in the training images. For color images this is 3
nc = 3

# Size of z latent vector (i.e. size of generator input)
nz = 100


# Size of feature maps in generator
ngf = 64

# Size of feature maps in discriminator
ndf = 64
```

```
CLASS torch.nn.ConvTranspose2d(in_channels, out_channels, kernel_size, stride=1, padding=0,
output_padding=0, groups=1, bias=True, dilation=1, padding_mode='zeros', device=None,
dtype=None) [SOURCE]
```

```
Generator(
  (main): Sequential(
    (0): ConvTranspose2d(100, 512, kernel_size=(4, 4), stride=(1, 1), bias=False)
    (1): BatchNorm2d(512, eps=1e-05, momentum=0.1, affine=True, track_running_stats=True)
    (2): ReLU(inplace=True)
    (3): ConvTranspose2d(512, 256, kernel_size=(4, 4), stride=(2, 2), padding=(1, 1), bias=False)
    (4): BatchNorm2d(256, eps=1e-05, momentum=0.1, affine=True, track_running_stats=True)
    (5): ReLU(inplace=True)
    (6): ConvTranspose2d(256, 128, kernel_size=(4, 4), stride=(2, 2), padding=(1, 1), bias=False)
    (7): BatchNorm2d(128, eps=1e-05, momentum=0.1, affine=True, track_running_stats=True)
    (8): ReLU(inplace=True)
    (9): ConvTranspose2d(128, 64, kernel_size=(4, 4), stride=(2, 2), padding=(1, 1), bias=False)
    (10): BatchNorm2d(64, eps=1e-05, momentum=0.1, affine=True, track_running_stats=True)
    (11): ReLU(inplace=True)
    (12): ConvTranspose2d(64, 3, kernel_size=(4, 4), stride=(2, 2), padding=(1, 1), bias=False)
    (13): Tanh()
  )
)
```

1 DCGAN



Discriminator Code

In [49]:

```
class Discriminator(nn.Module):
    def __init__(self, ngpu):
        super(Discriminator, self).__init__()
        self.ngpu = ngpu
        self.main = nn.Sequential(
            # input is `(nc) x 64 x 64`
            nn.Conv2d(nc, ndf, 4, 2, 1, bias=False),
            nn.LeakyReLU(0.2, inplace=True),
            # state size, `(ndf) x 32 x 32`
            nn.Conv2d(ndf, ndf * 2, 4, 2, 1, bias=False),
            nn.BatchNorm2d(ndf * 2),
            nn.LeakyReLU(0.2, inplace=True),
            # state size, `(ndf*2) x 16 x 16`
            nn.Conv2d(ndf * 2, ndf * 4, 4, 2, 1, bias=False),
            nn.BatchNorm2d(ndf * 4),
            nn.LeakyReLU(0.2, inplace=True),
            # state size, `(ndf*4) x 8 x 8`
            nn.Conv2d(ndf * 4, ndf * 8, 4, 2, 1, bias=False),
            nn.BatchNorm2d(ndf * 8),
            nn.LeakyReLU(0.2, inplace=True),
            # state size, `(ndf*8) x 4 x 4`
            nn.Conv2d(ndf * 8, 1, 4, 1, 0, bias=False),
            nn.Sigmoid()
        )

    def forward(self, input):
        return self.main(input)
```

```
# Number of channels in the training images. For color images this is 3
nc = 3

# Size of z latent vector (i.e. size of generator input)
nz = 100

# Size of feature maps in generator
ngf = 64


# Size of feature maps in discriminator
ndf = 64
```

```
CLASS torch.nn.Conv2d(in_channels, out_channels, kernel_size, stride=1, padding=0, dilation=1,
groups=1, bias=True, padding_mode='zeros', device=None, dtype=None) [SOURCE]
```

```
Discriminator(
  (main): Sequential(
    (0): Conv2d(3, 64, kernel_size=(4, 4), stride=(2, 2), padding=(1, 1), bias=False)
    (1): LeakyReLU(negative_slope=0.2, inplace=True)
    (2): Conv2d(64, 128, kernel_size=(4, 4), stride=(2, 2), padding=(1, 1), bias=False)
    (3): BatchNorm2d(128, eps=1e-05, momentum=0.1, affine=True, track_running_stats=True)
    (4): LeakyReLU(negative_slope=0.2, inplace=True)
    (5): Conv2d(128, 256, kernel_size=(4, 4), stride=(2, 2), padding=(1, 1), bias=False)
    (6): BatchNorm2d(256, eps=1e-05, momentum=0.1, affine=True, track_running_stats=True)
    (7): LeakyReLU(negative_slope=0.2, inplace=True)
    (8): Conv2d(256, 512, kernel_size=(4, 4), stride=(2, 2), padding=(1, 1), bias=False)
    (9): BatchNorm2d(512, eps=1e-05, momentum=0.1, affine=True, track_running_stats=True)
    (10): LeakyReLU(negative_slope=0.2, inplace=True)
    (11): Conv2d(512, 1, kernel_size=(4, 4), stride=(1, 1), bias=False)
    (12): Sigmoid()
  )
)
```

1 DCGAN

Loss Function



$$\min_G \max_D V(D, G) = \mathbb{E}_{\mathbf{x} \sim p_{\text{data}}(\mathbf{x})} [\log D(\mathbf{x})] + \mathbb{E}_{\mathbf{z} \sim p_{\mathbf{z}}(\mathbf{z})} [\log(1 - D(G(\mathbf{z})))]$$

Adversarial Training




Part 1 – Train the Discriminator

Part 2 – Train the Generator

1 DCGAN

Loss Function




$$\min_G \max_D V(D, G) = \mathbb{E}_{\mathbf{x} \sim p_{\text{data}}(\mathbf{x})} [\log D(\mathbf{x})] + \mathbb{E}_{\mathbf{z} \sim p_{\mathbf{z}}(\mathbf{z})} [\log(1 - D(G(\mathbf{z})))]$$

In practice, optimizing this generator objective does not work well!


Gradient signal dominated by region where sample is already good

When sample is likely fake, want to learn from it to improve generator. But gradient in this region is relatively flat!



1 DCGAN

Loss Function



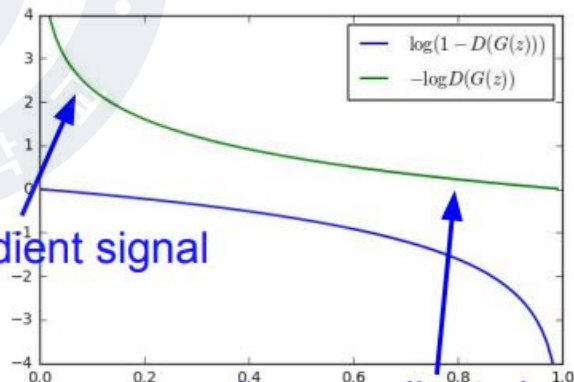
$$\min_G \max_D V(D, G) = \mathbb{E}_{\mathbf{x} \sim p_{\text{data}}(\mathbf{x})} [\log D(\mathbf{x})] + \mathbb{E}_{\mathbf{z} \sim p_{\mathbf{z}}(\mathbf{z})} [\log(1 - D(G(\mathbf{z})))]$$



$$\max_{\theta_g} \mathbb{E}_{\mathbf{z} \sim p(\mathbf{z})} \log(D_{\theta_d}(G_{\theta_g}(\mathbf{z})))$$

High gradient signal


Low gradient signal



Instead of minimizing likelihood of discriminator being correct, now maximize likelihood of discriminator being wrong. Same objective of fooling discriminator, but now higher gradient signal for bad samples => works much better!

1 DCGAN

Loss Function



$$\min_G \max_D V(D, G) = \mathbb{E}_{\mathbf{x} \sim p_{\text{data}}(\mathbf{x})} [\log D(\mathbf{x})] + \mathbb{E}_{\mathbf{z} \sim p_{\mathbf{z}}(\mathbf{z})} [\log(1 - D(G(\mathbf{z})))]$$

log maximization (=) negative log minimization



$$-(y_i \log(p_i) + (1 - y_i) \log(1 - p_i))$$


$$\min -1 \log D(\mathbf{x}) \quad \min - (1 - 0) \log(1 - D(G(\mathbf{z})))$$

```
critierion = nn.BCELoss()
# train D
loss = criterion(D(x), 1) + criterion(D(G(z)), 0)
loss.backward()
d_optimizer.step()

# train G
loss = criterion(D(G(z)), 1)
loss.backward()
g_optimizer.step()
```

1 DCGAN

Loss Function



$$\min_G \max_D V(D, G) = \mathbb{E}_{\mathbf{x} \sim p_{\text{data}}(\mathbf{x})} [\log D(\mathbf{x})] + \mathbb{E}_{\mathbf{z} \sim p_{\mathbf{z}}(\mathbf{z})} [\log(1 - D(G(\mathbf{z})))]$$

$$\max_{\theta_g} \mathbb{E}_{\mathbf{z} \sim p(\mathbf{z})} \log(D_{\theta_d}(G_{\theta_g}(\mathbf{z}))) \rightarrow \text{negative log minimization}$$



$$-(y_i \log(p_i) + (1 - y_i) \log(1 - p_i))$$

$$\min -1 \log D(G(z))$$

```
criterion = nn.BCELoss()
# train D
loss = criterion(D(x), 1) + criterion(D(G(z)), 0)
loss.backward()
d_optimizer.step()

# train G
loss = criterion(D(G(z)), 1)
loss.backward()
g_optimizer.step()
```

1 DCGAN

Loss Function Code

```
#####  
# (1) Update D network: maximize log(D(x)) + log(1 - D(G(z)))  
#####  
## Train with all-real batch  
netD.zero_grad()  
# Format batch  
real_cpu = data[0].to(device)  
b_size = real_cpu.size(0)  
label = torch.full((b_size,), real_label, dtype=torch.float, device=device)  
# Forward pass real batch through D  
output = netD(real_cpu).view(-1)  
# Calculate loss on all-real batch  
errD_real = criterion(output, label)  
# Calculate gradients for D in backward pass  
errD_real.backward()  
D_x = output.mean().item()
```

```
## Train with all-fake batch  
# Generate batch of latent vectors  
noise = torch.randn(b_size, nz, 1, 1, device=device)  
# Generate fake image batch with G  
fake = netG(noise)  
label.fill_(fake_label)  
# Classify all fake batch with D  
output = netD(fake.detach()).view(-1)  
# Calculate D's loss on the all-fake batch  
errD_fake = criterion(output, label)  
# Calculate the gradients for this batch, accumulated (summed) with previous gradients  
errD_fake.backward()  
D_G_z1 = output.mean().item()  
# Compute error of D as sum over the fake and the real batches  
errD = errD_real + errD_fake  
# Update D  
optimizerD.step()
```

In [53]:

```
# Initialize the 'BCELoss' function  
criterion = nn.BCELoss()  
  
# Create batch of latent vectors that we will use to visualize  
# the progression of the generator  
fixed_noise = torch.randn(64, nz, 1, 1, device=device)  
  
# Establish convention for real and fake labels during training  
real_label = 1.  
fake_label = 0.  
  
# Setup Adam optimizers for both G and D  
optimizerD = optim.Adam(netD.parameters(), lr=lr, betas=(beta1, 0.999))  
optimizerG = optim.Adam(netG.parameters(), lr=lr, betas=(beta1, 0.999))
```

$$\max_D V(D, G) = \mathbb{E}_{\mathbf{x} \sim p_{\text{data}}(\mathbf{x})} [\log D(\mathbf{x})] + \mathbb{E}_{\mathbf{z} \sim p_{\mathbf{z}}(\mathbf{z})} [\log(1 - D(G(\mathbf{z})))]$$

1 DCGAN

Loss Function Code

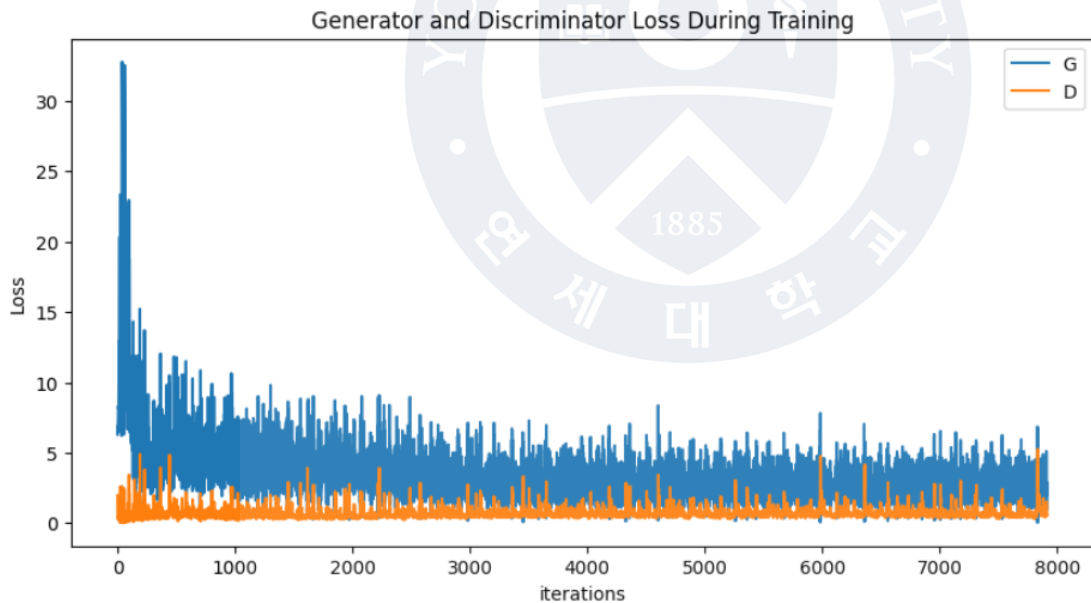
$$\max_{\theta_g} \mathbb{E}_{z \sim p(z)} \log(D_{\theta_d}(G_{\theta_g}(z)))$$

```
#####  
# (2) Update G network: maximize log(D(G(z)))  
#####  
netG.zero_grad()  
label.fill_(real_label) # fake labels are real for generator cost  
# Since we just updated D, perform another forward pass of all-fake batch through D  
output = netD(fake).view(-1)  
# Calculate G's loss based on this output  
errG = criterion(output, label)  
# Calculate gradients for G  
errG.backward()  
D_G_z2 = output.mean().item()  
# Update G  
optimizerG.step()
```

1 DCGAN



Loss During Training

```
In [55]: plt.figure(figsize=(10,5))
plt.title("Generator and Discriminator Loss During Training")
plt.plot(G_losses, label="G")
plt.plot(D_losses, label="D")
plt.xlabel("iterations")
plt.ylabel("Loss")
plt.legend()
plt.show()
```



1 DCGAN

Results



1 DCGAN

Github Code



ad-yonsei

Follow

DEEP LEARNING BASED ANOMALY
DETECTION MODELING (GEK6207.01-00),
Yonsei University

Block or Report

The screenshot shows a GitHub repository page for 'DCGAN' by user 'ad-yonsei'. The repository is on the 'main' branch, which is 1 commit ahead and 1 commit behind 'SkiddieAhm:main'. The file list includes 'data/celeba/img_align_celeba', 'README.md', and 'drgan_faces_tutorial.ipynb'. The repository description states it provides PyTorch tutorial code for understanding DCGAN (Deep Convolutional GAN) model, citing the original paper: 'Unsupervised Representation Learning with Deep Convolutional Generative Adversarial Networks' (ICLR 2016). A network pipeline diagram is shown at the bottom, illustrating the flow from a 100-dimensional latent vector 'z' through a 'Project and reshape' step to a 1024-dimensional layer, followed by two convolutional layers (CONV 1 and CONV 2) with various kernel sizes and stride 2 operations, leading to a final 3-dimensional output.

<https://github.com/ad-yonsei/Code-DCGAN>

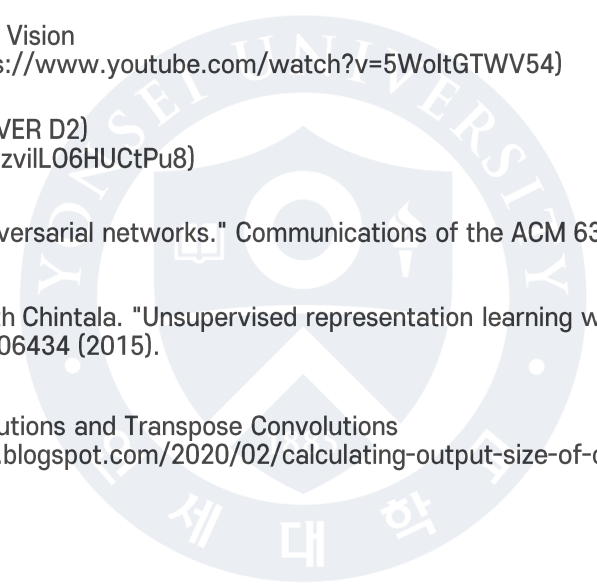
https://pytorch.org/tutorials/beginner/drgan_faces_tutorial.html

1 DCGAN



References

- ➡ CS231n: Deep Learning for Computer Vision
Lecture 13 | Generative Models (<https://www.youtube.com/watch?v=5WoltGTWV54>)
- ➡ Generative Adversarial Networks (NAVER D2)
(https://youtu.be/odpjk7_tGY0?si=d3zviiL06HUCtPu8)
- ➡ Goodfellow, Ian, et al. "Generative adversarial networks." Communications of the ACM 63.11 (2020): 139-144.
- ➡ Radford, Alec, Luke Metz, and Soumith Chintala. "Unsupervised representation learning with deep convolutional generative adversarial networks." arXiv preprint arXiv:1511.06434 (2015).
- ➡ Calculating the Output Size of Convolutions and Transpose Convolutions
(<https://makeyourownneuralnetwork.blogspot.com/2020/02/calculating-output-size-of-convolutions.html>)





Thank You

DCGAN - coding practice

Sunghyun Ahn

skd@yonsei.ac.kr

<2023/11/09>