

ECE695DL: Homework 7

Qilei Zhang

Due Date: Monday, Apr 11, 2022

1 Introduction

The goal of this homework is to understand the usage of generative adversarial network (GAN). GANs are a framework for teaching a DL model to capture the training data's distribution so model can generate new data from that same distribution. They are made of two distinct models, a generator and a discriminator. Particularly, generator code will use *Transpose Convolutions* to expand noise vectors into images that will look very similar to the images in the training data. The homework will use CelebFaces Attributes Dataset that contains thousands of celebrity images. The expected objective of the homework is to produce images will look very similar to the training images without being exactly the same as any of them.

2 Methodology

2.1 About Preparation

The basic information for the program shall be entered in the `FakeArg` class. `self.img_path` stores the downloaded images for training and the `batch_size` is the batch size settings. Here is the example scripts:

```
1 class FakeArg:
2     def __init__(self):
3         self.img_path = './ECE695/hw7/CelebA/Data/Train/'
4         self.batch_size = 24
```

`Dataloader` reads the images in the root folder. It will determine whether the set batch size will produce a remainder for the total number of training images. It will return the transformed normalized image tensor. Validation will also use the same `Dataloader` and `FakeArg`. Only the `FakeArg.img_pth` need to change to test image folder.

2.2 About Net Design

Two networks exists in the model, i.e., discriminator and generator. Discriminator uses the bases from previous homework. Here, discriminator takes a 3x64x64 input image. It

contains two convolution layers and two fully-connected layers. It also applied a sigmoid activation function before returning. Discriminator is a binary classification network that takes an image as input and outputs a scalar probability that the input image is real (as opposed to fake). The code of generator is from Professor Kak's `GeneratorDG1` Net in `AdversarialLearning`. Generator takes a vector with the length `nz`. It has five transpose convolutional layers and four batch normalization layers. It will return the output in the range of $[-1, 1]$ after applying `tanh` activation.

2.3 About GAN

The training of GAN network has two parts, which are discriminator training and generator training. First, the goal of training the discriminator is to maximize the probability of correctly classifying a given input as real or fake. Specifically, it want to maximize

$$\log(D(x)) + \log(1 - D(G(z)))$$

A batch of real samples from the training set calculate the loss $\log(D(x))$. A batch of fake samples computes the second term $\log(1 - D(G(z)))$. On the other hand, the goal of the generator is to minimize $\log(1 - D(G(z)))$ in order to generate better fakes. In summary the loss function is

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

It will use the Binary Cross Entropy loss (`BCELoss`) function as a criterion.

3 Implementation and Results

3.1 Main Program Code

3.1.1 hw07_training.py

```
1 import glob
2 import torch
3 import torchvision
4 import torch.utils.data
5
6 import imageio # For gif generation
7 from PIL import Image
8 from torch.utils.data import DataLoader, Dataset
9
10 import numpy as np
11 import torch.nn as nn
12 import torch.optim as optim
13 import torch.nn.functional as functional
14 import torchvision.transforms as tvt
15 import torchvision.transforms.functional as tvtF
```

```

16 import matplotlib.pyplot as plt
17
18
19 class FakeArg:
20     """
21     Put the fundamental information
22     """
23
24     def __init__(self):
25         self.img_path = '/Users/xxx/CelebA/Data/Train/'
26         self.batch_size = 24
27
28
29 class DatasetClass(Dataset):
30     def __init__(self, root, batch, trans=None):
31         self.root = root
32         self.img_info = []
33         self.batch = batch
34         self.transform = trans
35
36         search = self.root + "*"
37         for img in glob.glob(search):
38             self.img_info.append(img)
39
40         if len(self.img_info) % self.batch != 0:
41             remain = len(self.img_info) % self.batch
42             self.img_info = self.img_info[:-remain]
43
44     def __len__(self):
45         return len(self.img_info)
46
47     def __getitem__(self, idx):
48         img_path = self.img_info[idx]
49         image = Image.open(img_path)
50         im_ts = self.transform(image)
51         return im_ts
52
53
54 class DisNet(nn.Module):
55     def __init__(self):
56         super(DisNet, self).__init__()
57         self.conv1 = nn.Conv2d(3, 128, 3, padding=1)
58         self.conv2 = nn.Conv2d(128, 128, 3)
59         self.pool = nn.MaxPool2d(2, 2)
60         self.fc1 = nn.Linear(128 * 15 * 15, 1000)
61         self.fc2 = nn.Linear(1000, 1)

```

```

62     self.sig = nn.Sigmoid()
63
64     def forward(self, x):
65         x = self.pool(functional.relu(self.conv1(x)))
66         x = self.pool(functional.relu(self.conv2(x)))
67         x = x.view(-1, 128 * 15 * 15)
68         x = functional.relu(self.fc1(x))
69         x = self.fc2(x)
70         x = self.sig(x)
71         return x
72
73
74 # This network structure is referred from Prof. Kak's GeneratorDG1 Net.
75 class GenNet(nn.Module):
76     def __init__(self):
77         super(GenNet, self).__init__()
78         self.latent_to_image = nn.ConvTranspose2d(200, 512, kernel_size=4,
79             ↪ stride=1, padding=0, bias=False)
80         self.upsampler2 = nn.ConvTranspose2d(512, 256, kernel_size=4,
81             ↪ stride=2, padding=1, bias=False)
82         self.upsampler3 = nn.ConvTranspose2d(256, 128, kernel_size=4,
83             ↪ stride=2, padding=1, bias=False)
84         self.upsampler4 = nn.ConvTranspose2d(128, 64, kernel_size=4,
85             ↪ stride=2, padding=1, bias=False)
86         self.upsampler5 = nn.ConvTranspose2d(64, 3, kernel_size=4,
87             ↪ stride=2, padding=1, bias=False)
88         self.bn1 = nn.BatchNorm2d(512)
89         self.bn2 = nn.BatchNorm2d(256)
90         self.bn3 = nn.BatchNorm2d(128)
91         self.bn4 = nn.BatchNorm2d(64)
92         self.tanh = nn.Tanh()
93
94     def forward(self, x):
95         x = self.latent_to_image(x)
96         x = torch.nn.functional.relu(self.bn1(x))
97         x = self.upsampler2(x)
98         x = torch.nn.functional.relu(self.bn2(x))
99         x = self.upsampler3(x)
100        x = torch.nn.functional.relu(self.bn3(x))
101        x = self.upsampler4(x)
102        x = torch.nn.functional.relu(self.bn4(x))
103        x = self.upsampler5(x)
104        x = self.tanh(x)
105        return x

```

```

103 # The structure of this code is taken from Professor Kak's
    ↪ AdversarialLearning.
104 def run_code_for_training(dis_net, gen_net, trainLoader, bs,
    ↪ learning_rate=1e-5,
105                               momentum_set=0.9, epochs=20, device='cuda:0'):
106     beta1 = 0.5 # ?
107     nz = 200 # Change
108     iteration_count = 0
109     dnet = dis_net.to(device)
110     gnet = gen_net.to(device)
111     fixed_noise = torch.randn(bs, nz, 1, 1, device=device)
112     real_label = 1
113     fake_label = 0
114     criterion = nn.BCELoss()
115     dis_optimizer = optim.Adam(dnet.parameters(), lr=learning_rate,
    ↪ betas=(beta1, 0.999))
116     gen_optimizer = optim.Adam(gnet.parameters(), lr=learning_rate,
    ↪ betas=(beta1, 0.999))
117
118     dloss_record = []
119     gloss_record = []
120     fake_img_list = []
121
122     for epoch in range(epochs):
123         g_losses_per_print_cycle = []
124         d_losses_per_print_cycle = []
125         for i, data in enumerate(trainLoader):
126             if i % 10 == 0:
127                 print(i)
128                 dnet.zero_grad() # ?
129                 real_img_inputs = data
130                 real_img_inputs = real_img_inputs.to(device)
131
132                 # Loss for real
133                 label = torch.full((bs,), real_label, dtype=torch.float,
    ↪ device=device)
134                 output = dnet(real_img_inputs).view(-1)
135                 dis_loss_for_reals = criterion(output, label)
136                 dis_loss_for_reals.backward()
137
138                 # max operation
139                 noise = torch.randn(bs, nz, 1, 1, device=device)
140                 fakes = gnet(noise)
141                 label.fill_(fake_label)
142                 output = dnet(fakes.detach()).view(-1)
143                 dis_loss_for_fakes = criterion(output, label)

```

```

144     dis_loss_for_fakes.backward()
145     dis_loss = dis_loss_for_reals + dis_loss_for_fakes
146     d_losses_per_print_cycle.append(dis_loss)
147     dis_optimizer.step()
148
149     # max-min optimization
150     gnet.zero_grad()
151     label.fill_(real_label)
152     output = dnet(fakes).view(-1)
153     gen_loss = criterion(output, label)
154     g_losses_per_print_cycle.append(gen_loss)
155     gen_loss.backward()
156     gen_optimizer.step()
157
158     dloss_record.append(dis_loss.item())
159     gloss_record.append(gen_loss.item())
160
161     if (i + 1) % 100 == 0:
162         mean_d_loss =
163             ↪ torch.mean(torch.FloatTensor(d_losses_per_print_cycle))
164         mean_g_loss =
165             ↪ torch.mean(torch.FloatTensor(g_losses_per_print_cycle))
166         print("\n[epoch:%d, batch:%5d] mean_D_loss=%7.4f
167             ↪ mean_G_loss=%7.4f" %
168             (epoch + 1, i + 1, mean_d_loss.item(),
169             ↪ mean_g_loss.item()))
170         d_losses_per_print_cycle = []
171         g_losses_per_print_cycle = []
172
173     if (iteration_count % 200 == 0) or ((epoch == epochs - 1) and (i
174     ↪ == len(trainLoader) - 1)):
175         with torch.no_grad():
176             fake = gnet(fixed_noise).detach().cpu()
177             fake_img_list.append(torchvision.utils.make_grid(fake,
178             ↪ padding=1, pad_value=1, normalize=True))
179         iteration_count += 1
180
181     dis_net_path = './dis_net.pth'
182     gen_net_path = './gen_net.pth'
183     torch.save(dnet.state_dict(), dis_net_path)
184     torch.save(gnet.state_dict(), gen_net_path)
185     print('Finished Training')
186
187     return dloss_record, gloss_record, fake_img_list

```

```

184 def plot_trained_fake_img(img_list, train_dataloader, device):
185     images_list = []
186     for fake_img in img_list:
187         image = tvtf.to_pil_image(fake_img)
188         image_array = np.array(image)
189         images_list.append(image_array)
190     imageio.mimsave("generation_animation.gif", images_list, fps=5)
191
192     plt.subplot(1, 1, 1)
193     plt.axis("off")
194     plt.title("Fake Images")
195     plt.imshow(np.transpose(img_list[-1], (1, 2, 0)))
196     plt.savefig("real_vs_fake_images.png")
197     plt.show()
198
199
200 if __name__ == '__main__':
201     # Training on GPU or CPU
202     if torch.cuda.is_available():
203         device = 'cuda:0'
204     else:
205         device = 'cpu'
206
207     # Basic Information
208     args = FakeArg()
209
210     # Load and normalize data
211     transform = tvf.Compose([tvf.ToTensor(), tvf.Normalize((0.5, 0.5, 0.5),
212     ↪ (0.5, 0.5, 0.5))])
213
214     trainSet = DatasetClass(args.img_path, args.batch_size, transform)
215     trainLoader = DataLoader(dataset=trainSet, batch_size=args.batch_size,
216     ↪ shuffle=True, num_workers=0)
217     print('Loader Created')
218     print(len(trainLoader), 'images')
219
220     discriminator = DisNet()
221     generator = GenNet()
222     print('Net Created')
223
224     DNetLoss, GNetLOSS, FakeImg = run_code_for_training(discriminator,
225     ↪ generator, trainLoader, args.batch_size,
226     ↪ learning_rate=1e-4,
227     ↪ momentum_set=0.9,
228     ↪ epochs=50,
229     ↪ device=device)

```

```

224
225     plt.figure(figsize=(10, 5))
226     plt.title("Generator and Discriminator Loss During Training")
227     plt.plot(DNetLoss, label="D")
228     plt.plot(GNetLOSS, label="G")
229     plt.xlabel("iterations")
230     plt.ylabel("Loss")
231     plt.legend()
232     plt.savefig("loss_training.png")
233     plt.show()
234
235     plot_trained_fake_img(FakeImg, trainLoader, device)
236
237     print('Done')

```

3.1.2 hw07_validation.py

```

1  import copy
2  import torch
3  import torchvision
4  import torch.nn as nn
5  import matplotlib.pyplot as plt
6  import torchvision.transforms as tvf
7  from torch.utils.data import DataLoader
8  from hw07_training import FakeArg, DatasetClass, DisNet, GenNet,
   → plot_trained_fake_img
9
10
11 def validation(dis_net, gen_net, val_loader, bs, device="cpu"):
12     dis_net = copy.deepcopy(dis_net)
13     gen_net = copy.deepcopy(gen_net)
14     dnet = dis_net.to(device)
15     gnet = gen_net.to(device)
16     nz = 200
17     fixed_noise = torch.randn(bs, nz, 1, 1, device=device)
18     real_label = 1
19     fake_label = 0
20     criterion = nn.BCELoss()
21
22     fake_img_list = []
23     dloss_record = []
24     gloss_record = []
25     for i, data in enumerate(val_loader):
26         print(i)
27         real_img_inputs = data
28         real_img_inputs = real_img_inputs.to(device)

```



```

29
30     # Loss for real
31     label = torch.full((bs,), real_label, dtype=torch.float,
32     ↪ device=device)
33     output = dnet(real_img_inputs).view(-1)
34     dis_loss_for_reals = criterion(output, label)
35
36     # max operation
37     noise = torch.randn(bs, nz, 1, 1, device=device)
38     fakes = gnet(noise)
39     label.fill_(fake_label)
40     output = dnet(fakes.detach()).view(-1)
41     dis_loss_for_fakes = criterion(output, label)
42     dis_loss = dis_loss_for_reals + dis_loss_for_fakes
43
44     # max-min optimization
45     label.fill_(real_label)
46     output = dnet(fakes).view(-1)
47     gen_loss = criterion(output, label)
48
49     dloss_record.append(dis_loss.item())
50     gloss_record.append(gen_loss.item())
51
52     if (i % 200 == 0) or (i == len(val_loader) - 1):
53         with torch.no_grad():
54             fake = gnet(fixed_noise).detach().cpu()
55             fake_img_list.append(torchvision.utils.make_grid(fake,
56             ↪ padding=1, pad_value=1, normalize=True))
57
58     return dloss_record, gloss_record, fake_img_list
59
60 if __name__ == '__main__':
61     print('start')
62     if torch.cuda.is_available():
63         device = 'cuda:0'
64     else:
65         device = 'cpu'
66
67     # Basic Information
68     args = FakeArg()
69     args.img_path = '/Users/XXXX/CelebA/Data/Test/'
70
71     transform = tvn.Compose([tvn.ToTensor(), tvn.Normalize((0.5, 0.5, 0.5),
72     ↪ (0.5, 0.5, 0.5))])

```

```

72     valSet = DatasetClass(args.img_path, args.batch_size, transform)
73     valLoader = DataLoader(dataset=valSet, batch_size=args.batch_size,
    ↪     shuffle=False, num_workers=0)
74     print('Loader Created')
75
76     discriminator = DisNet()
77     generator = GenNet()
78     print('Net Created')
79
80     discriminator.load_state_dict(torch.load("dis_net.pth",
    ↪     map_location=torch.device(device)))
81     discriminator.eval()
82     generator.load_state_dict(torch.load("gen_net.pth",
    ↪     map_location=torch.device(device)))
83     generator.eval()
84     DNetLoss, GNetLOSS, FakeImg = validation(discriminator, generator,
    ↪     valLoader, args.batch_size, device=device)
85
86     print('Validation Finished, Starting Plots')
87
88     plt.figure(figsize=(10, 5))
89     plt.title("Generator and Discriminator Loss During Testing")
90     plt.plot(DNetLoss, label="D")
91     plt.plot(GNetLOSS, label="G")
92     plt.xlabel("iterations")
93     plt.ylabel("Loss")
94     plt.legend()
95     plt.savefig("loss_testing.png")
96     plt.show()
97
98     plot_trained_fake_img(FakeImg, valLoader, device)
99
100    print('Done')

```

3.2 Results

The output are shown in following pages.

3.2.1 Training loss




Figure 1: Training loss with learning rate $1e-4$ and 50 epochs.

3.2.2 Validation loss




Figure 2: Validation Loss.

3.2.3 Real Images




Figure 3: Real Images.

3.2.4 Fake Training Images



Figure 4: Fake Training Images Output.

3.2.5 Fake Test Images




Figure 5: Fake Test Images Output.

4 Lessons Learned

1. It is expected that at the end of training (very large epochs) generator can always fool discriminator. However, when I tried a very large epochs, sometimes it not works as the small epochs, i.e., after some moment generator starts to diverge. Probably, the problem is that the discriminator overfits.
2. Google Drive operations can time out when the number of files or subfolders in a folder grows too large. This may bring additional process time to let dataloader read the images. The solution is to move files contained in one folder into sub-folders.
3. Two neural networks contest with each other in a game (in the form of a zero-sum game, where one agent's gain is another agent's loss).

5 Suggested Enhancements

1. Add checkpoints to ensure the unexpected situation such as network interruption. In that way, training could continue from the break point. Otherwise, the training needs to start again.
2. Try to train a network fed with specific labels, and generate a image with the appointed label.
3. Try training on higher pixel photos.