

Exercise week 13

Generative Adversarial Networks

In this programming exercise, we are going to implement DCGAN (*Deep Convolutional Generative Adversarial Networks*) introduced in (Radford et al., 2016). The implementation will be in PyTorch, and a similar tutorial, which this one is heavily influenced by, is available at https://pytorch.org/tutorials/beginner/dcgan_faces_tutorial.html.

- The goal is to implement a standalone python program that can be used for training and generating images.
- This guide will show you the essentials of the implementation. For a more complete version, see the accompanying solution proposal.
- We will train on the *Labeled Faces in the Wild* dataset, available from <http://vis-www.cs.umass.edu/lfw/>. This is a relatively small dataset, and fits well for this purpose.
- The implementation will run on both CPU and GPU, but is not implemented to make use of multiple GPUs.
- The exercise was developed on, and runs with Python 3.6 and PyTorch 1.1, but have not been tested on other versions.
- Feel free to experiment with different approaches, this implementation is just a suggestion, and not necessarily the best.

1 Import data

Use the function in listing 1 to download the dataset.

```
1 import subprocess
2
3 def maybe_download_lfw(download_root_dir):
4     download_dir = download_root_dir.joinpath('lfw')
5     if (
6         download_dir.exists() and
7         download_dir.is_dir() and
8         list(download_dir.iterdir())
9     ):
10        print("LFW dataset already downloaded")
11        return download_dir
12
13    tar_path = download_root_dir.joinpath('lfw.tgz')
14    if not tar_path.exists():
15        url = 'http://vis-www.cs.umass.edu/lfw/lfw.tgz'
16        cmd = ['curl', url, '-o', str(tar_path)]
17        print("Downloading LFW from {} to {}".format(url, tar_path))
18        subprocess.call(cmd)
19
20    if not download_dir.exists():
21        download_dir.mkdir()
22
23    print("Unpacking")
24    cmd = ['tar', 'xzf', str(tar_path), '-C', str(download_dir)]
25    subprocess.call(cmd)
26
27    return download_dir
```

Listing 1: Download data

Create a `data_loader` that we can use to iterate through the training dataset. We will resize the images to a shape of $(c, h, w) = (3, 64, 64)$, and normalize their values to a range of $[-1, 1]$ in accordance with the original paper (the input image values are in the range $[0, 1]$). As in the original paper, we are going to use a batch-size of 128.

```

1 import Path
2 import torch
3 import torchvision
4
5 data_dir = Path('where you would like to keep your downloaded data')
6 download_dir = maybe_download_lfw(data_dir)
7 target_size = 64
8 batch_size = 128
9 data_loader = torch.utils.data.DataLoader(
10     torchvision.datasets.ImageFolder(
11         str(download_dir),
12         transform=torchvision.transforms.Compose(
13             [
14                 torchvision.transforms.Resize((target_size, target_size)),
15                 torchvision.transforms.ToTensor(),
16                 torchvision.transforms.Normalize(
17                     (0.5, 0.5, 0.5), (0.5, 0.5, 0.5)
18                 ),
19             ]
20         ),
21     ),
22     batch_size=batch_size,
23     shuffle=True,
24 )
25
26 # Example how to iterate through one epoch
27 for image_batch, label_batch in data_loader:
28     # Do something with the images and labels

```

Listing 2: Import data

2 Training

This section will describe the discriminator network and the generator network. It will then go on to describe how to train them together.

2.1 Network implementation

The discriminator is a network that (in our case) takes a $64 \times 64 \times 3$ input image, and outputs a number with a value in $(0, 1)$ that we will interpret as the probability that the input image is a real image. Listing 3 display one way to implement the discriminator network in PyTorch.

The generator is (in our case) a network that takes an one-dimensional vector of length `latent_length` and outputs an image with shape $64 \times 64 \times 3$. A suggestion for the generator network is implemented in listing 4.

```

1 class Discriminator(torch.nn.Module):
2     def __init__(self):
3         super(Discriminator, self).__init__()
4         self.leaky_relu = torch.nn.LeakyReLU(0.2)
5         # Input shape (c, h, w): (3, 64, 64)
6         self.conv_1 = torch.nn.Conv2d(
7             in_channels=3,
8             out_channels=64,
9             kernel_size=4,
10            stride=2,
11            padding=1,
12            bias=False,
13        )
14        # Input shape (c, h, w): (64, 32, 32)
15        self.conv_2 = torch.nn.Conv2d(
16            in_channels=64,
17            out_channels=128,
18            kernel_size=4,
19            stride=2,
20            padding=1,
21            bias=False,
22        )
23        self.bn_1 = torch.nn.BatchNorm2d(128)
24        # Input shape (c, h, w): (128, 16, 16)
25        self.conv_3 = torch.nn.Conv2d(
26            in_channels=128,
27            out_channels=256,
28            kernel_size=4,
29            stride=2,
30            padding=1,
31            bias=False,
32        )
33        self.bn_2 = torch.nn.BatchNorm2d(256)
34        # Input shape (c, h, w): (256, 8, 8)
35        self.conv_4 = torch.nn.Conv2d(
36            in_channels=256,
37            out_channels=512,
38            kernel_size=4,
39            stride=2,
40            padding=1,
41            bias=False,
42        )
43        self.bn_3 = torch.nn.BatchNorm2d(512)
44        # Input shape (c, h, w): (512, 4, 4)
45        self.conv_5 = torch.nn.Conv2d(
46            in_channels=512,
47            out_channels=1,
48            kernel_size=4,
49            stride=1,
50            padding=0,
51            bias=False,
52        )
53
54        def forward(self, x):
55            x = self.leaky_relu(self.conv_1(x))
56            x = self.leaky_relu(self.bn_1(self.conv_2(x)))
57            x = self.leaky_relu(self.bn_2(self.conv_3(x)))
58            x = self.leaky_relu(self.bn_3(self.conv_4(x)))
59            x = torch.sigmoid(self.conv_5(x))
60            return x

```

Listing 3: Discriminator network

```

1 latent_length = 100
2 class Generator(torch.nn.Module):
3     def __init__(self):
4         super(Generator, self).__init__()
5         # Input shape (c, h, w): (latent_length, 1, 1)
6         self.convtr_1 = torch.nn.ConvTranspose2d(
7             in_channels=latent_length,
8             out_channels=512,
9             kernel_size=4,
10            stride=1,
11            padding=0,
12            bias=False,
13        )
14        self.bn_1 = torch.nn.BatchNorm2d(num_features=512)
15        # Input shape (c, h, w): (512, 4, 4)
16        self.convtr_2 = torch.nn.ConvTranspose2d(
17            in_channels=512,
18            out_channels=256,
19            kernel_size=4,
20            stride=2,
21            padding=1,
22            bias=False,
23        )
24        self.bn_2 = torch.nn.BatchNorm2d(num_features=256)
25        # Input shape (c, h, w): (256, 8, 8)
26        self.convtr_3 = torch.nn.ConvTranspose2d(
27            in_channels=256,
28            out_channels=128,
29            kernel_size=4,
30            stride=2,
31            padding=1,
32            bias=False,
33        )
34        self.bn_3 = torch.nn.BatchNorm2d(num_features=128)
35        # Input shape (c, h, w): (128, 16, 16)
36        self.convtr_4 = torch.nn.ConvTranspose2d(
37            in_channels=128,
38            out_channels=64,
39            kernel_size=4,
40            stride=2,
41            padding=1,
42            bias=False,
43        )
44        self.bn_4 = torch.nn.BatchNorm2d(num_features=64)
45        # Input shape (c, h, w): (64, 32, 32)
46        self.convtr_5 = torch.nn.ConvTranspose2d(
47            in_channels=64,
48            out_channels=3,
49            kernel_size=4,
50            stride=2,
51            padding=1,
52            bias=False,
53        )
54        # Output shape (c, h, w): (3, 64, 64). NOTE: The spatial shape should
55        # match target_size
56    def forward(self, x):
57        x = torch.nn.functional.relu(self.bn_1(self.convtr_1(x)))
58        x = torch.nn.functional.relu(self.bn_2(self.convtr_2(x)))
59        x = torch.nn.functional.relu(self.bn_3(self.convtr_3(x)))
60        x = torch.nn.functional.relu(self.bn_4(self.convtr_4(x)))
61        x = torch.tanh(self.convtr_5(x))
62        return x

```

Listing 4: Generator network

2.2 Parameter initialisation

All weights are initialised from a zero-centered Normal distribution with standard deviation 0.02, while the Batch Norm bias parameters are initialised to zero. This can be achieved in PyTorch by implementing the function

```
1 def initialise_weights(submodule):
2     if (
3         isinstance(submodule, torch.nn.Conv2d) or
4         isinstance(submodule, torch.nn.ConvTranspose2d)
5     ):
6         # Initialise from a random normal distribution with mean 0.0 and
7         # stdev 0.02
8         torch.nn.init.normal_(submodule.weight.data, 0.0, 0.02)
9     elif isinstance(submodule, torch.nn.BatchNorm2d):
10        # Initialise from a random normal distribution with mean 1.0 and
11        # stdev 0.02
12        torch.nn.init.normal_(submodule.weight.data, 1.0, 0.02)
13        # Initialise all bias parameters to zero
14        torch.nn.init.constant_(submodule.bias.data, 0.0)
```

Listing 5: Weight initialisation

and applying it on the discriminator network and the generator network

```
1 device = torch.device('cuda:0' if torch.cuda.is_available() else 'cpu')
2 discriminator = Discriminator().to(device)
3 discriminator.apply(initialise_weights)
4 generator = Generator().to(device)
5 generator.apply(initialise_weights)
```

Listing 6: Weight initialisation

2.3 Adversarial training

For each update step we are going to update the discriminator network D , and the generator network G once. First, consider the discriminator network with the associated loss

$$L_D = -\frac{1}{m} \sum_{i=1}^m [\log(D(x_i)) + \log(1 - D(G(z_i)))] \quad (2.1)$$

over a mini-batch of size m , as shown in the lecture slides for week 13. This is a sum of two binary-cross-entropy loss functions; one where the discriminator output on real images x_i is compared against 1, and one where the discriminator output on generated images $G(z_i)$ is compared against 0. For the discriminator network, we use the Adam optimisation method with the same hyperparameter values as in the DCGAN paper. The discriminator update can be implemented as.

```

1 binary_ce = torch.nn.BCELoss()
2 discriminator_optimiser = torch.optim.Adam(
3     discriminator.parameters(),
4     lr=0.0002,
5     betas=(0.5, 0.999),
6 )
7 for image_batch, _ in data_loader:
8     discriminator.zero_grad()
9     image_batch = image_batch.to(device)
10    batch_size = image_batch.shape[0]
11
12    # Discriminator wrt true images
13    real_label_batch = torch.full((batch_size, ), real_label, device=device)
14    discriminator_output_real = discriminator.forward(image_batch).view(-1)
15    discriminator_loss_real = binary_ce(
16        discriminator_output_real, real_label_batch
17    )
18    discriminator_loss_real.backward()
19
20    # Discriminator wrt generated images
21    noise = torch.randn(batch_size, latent_length, 1, 1, device=device)
22    fake_batch = generator.forward(noise)
23    fake_label_batch = torch.full((batch_size, ), fake_label, device=device)
24    discriminator_output_fake = discriminator\
25        .forward(fake_batch.detach())\
26        .view(-1)
27    discriminator_loss_fake = binary_ce(
28        discriminator_output_fake, fake_label_batch
29    )
30    discriminator_loss_fake.backward()
31
32    discriminator_optimiser.step()

```

Listing 7: Discriminator network update

Note that the binary cross entropy for a single example between a label $y \in \{0, 1\}$ and a prediction $\hat{y} \in [0, 1]$ is

$$l = -[y \log \hat{y} + (1 - y) \log(1 - \hat{y})].$$

Since, for the discriminator loss, the labels for real examples are 1 and for generated examples are 0, we end up with the expression in eq. (2.1). Also note the `fake_batch.detach()` on line 25 in the listing 7, this tells PyTorch not to compute gradients for the generator here.

The generator will be updated similarly. Note that we already have computed the forward pass of the generator (line 22 in listing 7). Since we have updated the discriminator, we evaluate it again on the generated images (line 8 in listing 8). From the lectures, we have that the generator loss is

$$L_G = -\frac{1}{m} \sum_{i=1}^m \log D(G(z_i)). \quad (2.2)$$

This is the binary cross entropy loss comparing generated images $G(z_i)$ with the label 1. Remember that the objective is to guide the generator to generate images that looks like they are from the same distribution as the real images.

```

1 generator_optimiser = torch.optim.Adam(
2     generator.parameters(),
3     lr=0.0002,
4     betas=(0.5, 0.999),
5 )
6 for image_batch, _ in data_loader:
7     generator.zero_grad()
8     discriminator_output_fake = discriminator.forward(fake_batch).view(-1)
9     generator_loss = binary_ce(discriminator_output_fake, real_label_batch)
10    generator_loss.backward()
11
12    generator_optimiser.step()

```

Listing 8: Generator network update

3 Restore a trained model

The point of this exercise is to train the above adversarial network so that we can use the trained generator to generate new examples from the training data distribution. For us to be able to restore the generator, we need to save its parameter values. This can be achieved by

```

1 torch.save(generator.state_dict(), "Path where you want to save your model")

```

Listing 9: Save generator parameter values

which needs to be written inside the for-loop in listing 7, and activated at appropriate times, for example at every n steps. The generator can then be restored with

```

1 noise = torch.randn(64, latent_length, 1, 1, device=device)
2 generator.load_state_dict(torch.load("Path with desired checkpoint"))
3 generator.eval()
4 generated_images = generator.forward(noise)

```

Listing 10: Restore generator

4 Example results

You can display the generated images with the code in listing 11.

```

1 def plot_mosaic(images, filename):
2     num_cols = 8
3     num_rows = min(images.shape[0] // num_cols, 8)
4     plt.figure()
5     images = images[:num_rows*num_cols, :, :, :]
6     images = images * 0.5 + 0.5 # Invert normalisation
7     image_grid = np.transpose(
8         torchvision.utils.make_grid(images, padding=2).cpu().detach().numpy(),
9         (1, 2, 0)
10    )
11    plt.axis('off')
12    plt.imshow(image_grid)
13 plot_mosaic(generated_images, "Location you want to put output images")

```

Listing 11: Restore generator




Figure 4.1: Generated images from generator restored after training 1 000 steps




Figure 4.2: Generated images from generator restored after training 2 000 steps




Figure 4.3: Generated images from generator restored after training 3 000 steps




Figure 4.4: Generated images from generator restored after training 4 000 steps




Figure 4.5: Generated images from generator restored after training 5 000 steps



Figure 4.6: Generated images from generator restored after training 6 000 steps



Figure 4.7: Generated images from generator restored after training 7 000 steps



Figure 4.8: Generated images from generator restored after training 8 000 steps



Figure 4.9: Generated images from generator restored after training 9 000 steps



Figure 4.10: Generated images from generator restored after training 10 000 steps

References

Alec Radford, Luke Metz, and Soumith Chintala. Unsupervised Representation Learning with Deep Convolutional Generative Adversarial Networks. In *ICLR 2016*, pages 1–16, 2016. URL <http://arxiv.org/abs/1511.06434>.