

# Demo of Convolutional Neural Network on CIFAR10

David I. Inouye

# Demo of Convolutional Neural Network in PyTorch on CIFAR10

Adapted from PyTorch tutorial from (skipping details):

[https://pytorch.org/tutorials/beginner/blitz/cifar10\\_tutorial.html](https://pytorch.org/tutorials/beginner/blitz/cifar10_tutorial.html)

# Load data (skipping details see tutorial for details)

```
1 import torch
2 import torchvision
3 import torchvision.transforms as transforms
4
5 transform = transforms.Compose(
6     [transforms.ToTensor(),
7      transforms.Normalize((0.5, 0.5, 0.5), (0.5, 0.5, 0.5))])
8
9 trainset = torchvision.datasets.CIFAR10(
10     root='./data', train=True,
11     download=True, transform=transform)
12 trainloader = torch.utils.data.DataLoader(
13     trainset, batch_size=4,
14     shuffle=True, num_workers=2)
15
16 testset = torchvision.datasets.CIFAR10(
17     root='./data', train=False,
18     download=True, transform=transform)
19 testloader = torch.utils.data.DataLoader(
20     testset, batch_size=4,
21     shuffle=False, num_workers=2)
22
23 classes = ('plane', 'car', 'bird', 'cat',
24            'deer', 'dog', 'frog', 'horse', 'ship', 'truck')
```

# Show some images from CIFAR10

```
1 import matplotlib.pyplot as plt
2 import numpy as np
3
4 # functions to show an image
5 def imshow(img):
6     img = img / 2 + 0.5     # unnormalize
7     npimg = img.numpy()
8     plt.imshow(np.transpose(npimg, (1, 2, 0)))
9     plt.show()
10
11 # get some random training images
12 dataiter = iter(trainloader)
13 images, labels = next(dataiter)
14
15 # show images
16 imshow(torchvision.utils.make_grid(images))
17 # print labels
18 print(' '.join('%5s' % classes[labels[j]] for j in range(4)))
```



horse horse truck frog

# Define a Convolutional Neural Network

```
1 import torch.nn as nn
2 import torch.nn.functional as F
3
4 class Net(nn.Module):
5     def __init__(self):
6         super(Net, self).__init__()
7         # nn.Conv2d(in_channels, out_channels, kernel_size)
8         self.conv1 = nn.Conv2d(3, 6, 5)
9         # nn.MaxPool2d(kernel_size, stride)
10        self.pool = nn.MaxPool2d(2, 2)
11        self.conv2 = nn.Conv2d(6, 16, 5)
12        # nn.Linear(in_features, out_features)
13        self.fc1 = nn.Linear(16 * 5 * 5, 120)
14        self.fc2 = nn.Linear(120, 84)
15        self.fc3 = nn.Linear(84, 10)
16
17    def forward(self, x):
18        # Input is (N, 3, 32, 32)
19        x = F.relu(self.conv1(x)) # (N, 6, 28, 28)
20        x = self.pool(x) # (N, 6, 14, 14)
21        x = F.relu(self.conv2(x)) # (N, 16, 10, 10)
22        x = self.pool(x) # (N, 16, 5, 5)
23        x = x.view(-1, 16 * 5 * 5) # (N, 400)
24        x = F.relu(self.fc1(x)) # (N, 120)
25        x = F.relu(self.fc2(x)) # (N, 84)
26        x = self.fc3(x) # (N, 10)
27        return x
```

# Parameters of layers

- `torch.nn.Conv2d` and similar functions produce object that automatically registers its parameters inside the `torch.nn.Module`
- Thus, when calling `model.parameters()`, it will include these parameters
- Note that simple ReLU and maxpool functions do not have parameters

```
1 # Remember convolution weight has size (out_channels, in_chan
2 total_num_params = 0
3 for name, p in net.named_parameters():
4     total_num_params += p.numel() # Number of elements
5     print(name, ',', p.size(), type(p))
6 print(f'Total number of parameters: {total_num_params}')
```

```
conv1.weight , torch.Size([6, 3, 5, 5]) <class
'torch.nn.parameter.Parameter'>
conv1.bias , torch.Size([6]) <class
'torch.nn.parameter.Parameter'>
conv2.weight , torch.Size([16, 6, 5, 5]) <class
'torch.nn.parameter.Parameter'>
conv2.bias , torch.Size([16]) <class
'torch.nn.parameter.Parameter'>
fc1.weight , torch.Size([120, 400]) <class
'torch.nn.parameter.Parameter'>
fc1.bias , torch.Size([120]) <class
'torch.nn.parameter.Parameter'>
fc2.weight , torch.Size([84, 120]) <class
'torch.nn.parameter.Parameter'>
fc2.bias , torch.Size([84]) <class
'torch.nn.parameter.Parameter'>
fc3.weight , torch.Size([10, 84]) <class
```

# Define a Loss function and optimizer

Let's use a Classification Cross-Entropy loss and SGD with momentum.

```
1 import torch.optim as optim
2
3 criterion = nn.CrossEntropyLoss()
4 optimizer = optim.SGD(
5     net.parameters(), lr=0.001, momentum=0.9)
```

# Train the network

This is when things start to get interesting. We simply have to loop over our data iterator, and feed the inputs to the network and optimize.

```
1 for epoch in range(2): # Loop over dataset
2     running_loss = 0.0
3     for i, data in enumerate(trainloader, 0):
4         inputs, labels = data
5         # zero the parameter gradients
6         optimizer.zero_grad()
7
8         # forward + backward + optimize
9         outputs = net(inputs)
10        loss = criterion(outputs, labels)
11        loss.backward()
12        optimizer.step()
13
14        # print statistics
15        running_loss += loss.item()
16        if i % 2000 == 1999: # print every 2000 mini-batches
17            print('[%d, %5d] loss: %.3f' %
18                  (epoch + 1, i + 1, running_loss / 2000))
19            running_loss = 0.0
20 print('Finished Training')
```

```
[1, 2000] loss: 2.238
[1, 4000] loss: 1.945
[1, 6000] loss: 1.730
[1, 8000] loss: 1.614
[1, 10000] loss: 1.555
[1, 12000] loss: 1.493
[2, 2000] loss: 1.432
[2, 4000] loss: 1.396
[2, 6000] loss: 1.377
[2, 8000] loss: 1.344
[2, 10000] loss: 1.323
[2, 12000] loss: 1.290
Finished Training
```

# Saving and loading trained model

- Note: If you want to restart training, you also need to save the optimizer states.


```
1 PATH = './cifar_net.pth'  
2 torch.save(net.state_dict(), PATH)  
3  
4 # Load the model from saved path  
5 net = Net()  
6 net.load_state_dict(torch.load(PATH))
```

```
<All keys matched successfully>
```

See [here <https://pytorch.org/docs/stable/notes/serialization.html>](https://pytorch.org/docs/stable/notes/serialization.html) for more details on saving PyTorch models.

# Check the network on the test data

```
1 dataiter = iter(testloader)
2 images, labels = next(dataiter)
3
4 # Show images with ground truth
5 imshow(torchvision.utils.make_grid(images))
6 print('GroundTruth: ', ' '.join('%5s' % classes[labels[j]] fo
7
8 # Predict outputs
9 outputs = net(images)
10 _, predicted = torch.max(outputs, 1)
11
12 print('Predicted: ', ' '.join('%5s' % classes[predicted[j]]
13       for j in range(4)))
```



```
GroundTruth:  cat  ship  ship plane
Predicted:   cat  car  car  ship
```

# Evaluate on test dataset

```
1 correct = 0
2 total = 0
3 with torch.no_grad():
4     for data in testloader:
5         images, labels = data
6         outputs = net(images)
7         _, predicted = torch.max(outputs.data, 1)
8         total += labels.size(0)
9         correct += (predicted == labels).sum().item()
10
11 print(
12     'Accuracy of the network on the 10000 test images: %d %%'
13     % (100 * correct / total)
14 )
```

Accuracy of the network on the 10000 test images: 55 %

- Note that 10% is random chance.

# Let's evaluate per-class accuracies to see which ones performed well

```
1 class_correct = list(0. for i in range(10))
2 class_total = list(0. for i in range(10))
3 with torch.no_grad():
4     for data in testloader:
5         images, labels = data
6         outputs = net(images)
7         _, predicted = torch.max(outputs, 1)
8         c = (predicted == labels).squeeze()
9         for i in range(4):
10            label = labels[i]
11            class_correct[label] += c[i].item()
12            class_total[label] += 1
13
14
15 for i in range(10):
16     print('Accuracy of %5s : %2d %%' % (
17         classes[i], 100 * class_correct[i] / class_total[i]))
```

```
Accuracy of plane : 70 %
Accuracy of  car : 71 %
Accuracy of  bird : 27 %
Accuracy of  cat : 27 %
Accuracy of  deer : 53 %
Accuracy of  dog : 52 %
Accuracy of  frog : 79 %
Accuracy of horse : 57 %
Accuracy of  ship : 56 %
Accuracy of truck : 55 %
```

# What about running on GPUs?

```
1 device = torch.device("cuda:0" if torch.cuda.is_available() else
2
3 # Assuming that we are on a CUDA machine, this should print a
4 print(device) # Would print GPU if available.
5
6 # This would move all model parameters to device`
7 net.to(device)
8
9 # You would also have to move inputs and target to device
10 data = next(iter(trainloader))
11 inputs, labels = data[0].to(device), data[1].to(device)
```

cpu