

PyTorch Tutorial from:

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)

```
In [2]: %matplotlib inline

import torch
import torchvision
import torchvision.transforms as transforms

transform = transforms.Compose(
    [transforms.ToTensor(),
     transforms.Normalize((0.5, 0.5, 0.5), (0.5, 0.5, 0.5))])

trainset = torchvision.datasets.CIFAR10(root='./data', train=True,
                                       download=True, transform=transform)
trainloader = torch.utils.data.DataLoader(trainset, batch_size=4,
                                          shuffle=True, num_workers=2)

testset = torchvision.datasets.CIFAR10(root='./data', train=False,
                                       download=True, transform=transform)
testloader = torch.utils.data.DataLoader(testset, batch_size=4,
                                         shuffle=False, num_workers=2)

classes = ('plane', 'car', 'bird', 'cat',
           'deer', 'dog', 'frog', 'horse', 'ship', 'truck')

import matplotlib.pyplot as plt
import numpy as np

# functions to show an image

def imshow(img):
    img = img / 2 + 0.5     # unnormalize
    npimg = img.numpy()
    plt.imshow(np.transpose(npimg, (1, 2, 0)))
    plt.show()

# get some random training images
dataiter = iter(trainloader)
images, labels = dataiter.next()

# show images
imshow(torchvision.utils.make_grid(images))
# print labels
print(' '.join('%5s' % classes[labels[j]] for j in range(4)))
```

Files already downloaded and verified
 Files already downloaded and verified



bird deer plane plane

Define a Convolutional Neural Network

```
In [3]: import torch.nn as nn
import torch.nn.functional as F

class Net(nn.Module):
    def __init__(self):
        super(Net, self).__init__()
        # nn.Conv2d(in_channels, out_channels/n_filters, kernel_size)
        self.conv1 = nn.Conv2d(3, 6, 5)
        # nn.MaxPool2d(kernel_size, stride)
        self.pool = nn.MaxPool2d(2, 2)
        self.conv2 = nn.Conv2d(6, 16, 5)
        # nn.Linear(in_features, out_features)
        self.fc1 = nn.Linear(16 * 5 * 5, 120)
        self.fc2 = nn.Linear(120, 84)
        self.fc3 = nn.Linear(84, 10)

    def forward(self, x):
        x = self.pool(F.relu(self.conv1(x)))
        x = self.pool(F.relu(self.conv2(x)))
        x = x.view(-1, 16 * 5 * 5)
        x = F.relu(self.fc1(x))
        x = F.relu(self.fc2(x))
        x = self.fc3(x)
        return x

net = Net()
```

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

```
In [22]: # Remember convolution weight has size (out_channels, in_channels, *kernel_size)
for name, p in net.named_parameters():
    print(name, ', ', p.size(), type(p))
    #print(type(p))
    #print(p)
```

```
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 'torch.nn.parameter.Parameter'>
fc3.bias , torch.Size([10]) <class 'torch.nn.parameter.Parameter'>
```

Define a Loss function and optimizer

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

```
In [6]: import torch.optim as optim

criterion = nn.CrossEntropyLoss()
optimizer = optim.SGD(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.

```
In [7]: for epoch in range(2): # loop over the dataset multiple times

        running_loss = 0.0
        for i, data in enumerate(trainloader, 0):
            # get the inputs; data is a list of [inputs, labels]
            inputs, labels = data

            # zero the parameter gradients
            optimizer.zero_grad()

            # forward + backward + optimize
            outputs = net(inputs)
            loss = criterion(outputs, labels)
            loss.backward()
            optimizer.step()

            # print statistics
            running_loss += loss.item()
            if i % 2000 == 1999: # print every 2000 mini-batches
                print('[%d, %5d] loss: %.3f' %
                      (epoch + 1, i + 1, running_loss / 2000))
                running_loss = 0.0

        print('Finished Training')
```

```
[1, 2000] loss: 2.154
[1, 4000] loss: 1.856
```

```
170500096it [00:40, 6815626.28it/s]
```

```
[1, 6000] loss: 1.670
[1, 8000] loss: 1.601
[1, 10000] loss: 1.533
[1, 12000] loss: 1.482
[2, 2000] loss: 1.417
[2, 4000] loss: 1.374
[2, 6000] loss: 1.365
[2, 8000] loss: 1.362
[2, 10000] loss: 1.356
[2, 12000] loss: 1.318
Finished Training
```

Let's quickly save our trained model:

```
In [8]: PATH = './cifar_net.pth'
torch.save(net.state_dict(), PATH)
```

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

Test the network on the test data


We have trained the network for 2 passes over the training dataset. But we need to check if the network has learnt anything at all.

We will check this by predicting the class label that the neural network outputs, and checking it against the ground-truth. If the prediction is correct, we add the sample to the list of correct predictions.

Okay, first step. Let us display an image from the test set to get familiar.

```
In [17]: dataiter = iter(testloader)
images, labels = dataiter.next()

# print images
imshow(torchvision.utils.make_grid(images))
print('GroundTruth: ', ' '.join('%5s' % classes[labels[j]] for j in range(4)))
```



GroundTruth: cat ship ship plane

Next, let's load back in our saved model (note: saving and re-loading the model wasn't necessary here, we only did it to illustrate how to do so):

```
In [18]: net = Net()
net.load_state_dict(torch.load(PATH))
```

```
Out[18]: IncompatibleKeys(missing_keys=[], unexpected_keys=[])
```

Okay, now let us see what the neural network thinks these examples above are:

```
In [19]: outputs = net(images)
```

The outputs are energies for the 10 classes. The higher the energy for a class, the more the network thinks that the image is of the particular class. So, let's get the index of the highest energy:

```
In [20]: _, predicted = torch.max(outputs, 1)
print('Predicted: ', ' '.join('%5s' % classes[predicted[j]]
                                for j in range(4)))
Predicted:   cat   car  ship  ship
```

The results seem pretty good.

Let us look at how the network performs on the whole dataset.

```
In [21]: correct = 0
total = 0
with torch.no_grad():
    for data in testloader:
        images, labels = data
        outputs = net(images)
        _, predicted = torch.max(outputs.data, 1)
        total += labels.size(0)
        correct += (predicted == labels).sum().item()

print('Accuracy of the network on the 10000 test images: %d %%' % (
    100 * correct / total))
Accuracy of the network on the 10000 test images: 53 %
```

That looks way better than chance, which is 10% accuracy (randomly picking a class out of 10 classes). Seems like the network learnt something.

Hmmm, what are the classes that performed well, and the classes that did not perform well:

```
In [22]: class_correct = list(0. for i in range(10))
class_total = list(0. for i in range(10))
with torch.no_grad():
    for data in testloader:
        images, labels = data
        outputs = net(images)
        _, predicted = torch.max(outputs, 1)
        c = (predicted == labels).squeeze()
        for i in range(4):
            label = labels[i]
            class_correct[label] += c[i].item()
            class_total[label] += 1

for i in range(10):
    print('Accuracy of %5s : %2d %%' % (
        classes[i], 100 * class_correct[i] / class_total[i]))
```

```
Accuracy of plane : 53 %
Accuracy of car : 84 %
Accuracy of bird : 44 %
Accuracy of cat : 35 %
Accuracy of deer : 28 %
Accuracy of dog : 44 %
Accuracy of frog : 63 %
Accuracy of horse : 60 %
Accuracy of ship : 59 %
Accuracy of truck : 60 %
```

Okay, so what next?

How do we run these neural networks on the GPU?

Training on GPU

Just like how you transfer a Tensor onto the GPU, you transfer the neural net onto the GPU.

Let's first define our device as the first visible cuda device if we have CUDA available:

```
In [15]: device = torch.device("cuda:0" if torch.cuda.is_available() else "cpu")

# Assuming that we are on a CUDA machine, this should print a CUDA device:

print(device)
```

```
cpu
```


The rest of this section assumes that `device` is a CUDA device.

Then these methods will recursively go over all modules and convert their parameters and buffers to CUDA tensors:

```
.. code:: python

    net.to(device)
```

Remember that you will have to send the inputs and targets at every step to the GPU too:

```
.. code:: python

    inputs, labels = data[0].to(device), data[1].to(device)
```

Why dont I notice MASSIVE speedup compared to CPU? Because your network is really small.

Exercise: Try increasing the width of your network (argument 2 of the first `nn.Conv2d`, and argument 1 of the second `nn.Conv2d` – they need to be the same number), see what kind of speedup you get.

Goals achieved:

- Understanding PyTorch's Tensor library and neural networks at a high level.
- Train a small neural network to classify images

Training on multiple GPUs

If you want to see even more MASSIVE speedup using all of your GPUs, please check out [:doc: data_parallel_tutorial](#).

Where do I go next?

- [:doc: Train neural nets to play video games](#)
- [Train a state-of-the-art ResNet network on imagenet](#)
- [Train a face generator using Generative Adversarial Networks](#)
- [Train a word-level language model using Recurrent LSTM networks](#)
- [More examples](#)
- [More tutorials](#)
- [Discuss PyTorch on the Forums](#)
- [Chat with other users on Slack](#)

In []: