
pytorch-nlp-tutorial-sf2017

Documentation

Release

Brian McMahan and Delip Rao

May 01, 2018

Extra Resources

1	Getting the Data	1
1.1	Option 1: Download and Setup things on your laptop	1
1.2	Option 2: Use O'Reilly's online resource through your browser	1
2	Environment Setup	3
2.1	0. Get Anaconda	3
2.2	1. Create a new environment	3
2.3	2. Install Dependencies	4
3	Frequency Asked Questions	7
3.1	Do I Need to have a NVIDIA GPU enabled laptop?	7
4	Migrating to PyTorch 0.4.0	9
5	Solutions	11
5.1	Problem 1	11
5.2	Problem 2	11
6	Warm Up Exercise	13
7	Fail Fast Prototype Mode	15
7.1	Prototyping an embedding	15
8	Tensor-Fu-1	17
8.1	Exercise 1	17
8.2	Exercise 2	17
8.3	Exercise 3	17
9	Tensor-Fu-2	19
9.1	Exercise 1	19
9.2	Exercise 2	19
10	Exercise: Interpolating Between Vectors	21
11	Exercise: Sampling from an RNN	23
12	Design Pattern: Attention	25

13 General Information	27
13.1 Prerequisites:	27
13.2 Hardware and/or installation requirements:	27

In this training, there are two options of participating.

1.1 Option 1: Download and Setup things on your laptop

The first option is to download the data below, setup the environment, and download the notebooks when we make them available. If you choose this options but do not download the data before the first day, we will have several flash drives with the data on it.

Please visit [this link](#) to download the data.

1.2 Option 2: Use O'Reilly's online resource through your browser

The second option is to use an online resource provided by O'Reilly. On the first day of this training, you will be provided with a link to a JupyterHub instance where the environment will be pre-made and ready to go! If you choose this option, you do not have to do anything until you arrive on Sunday. You are still required to bring your laptop.

Environment Setup

On this page, you will find not only the list of dependencies to install for the tutorial, but a description of how to install them. This tutorial assumes you have a laptop with OSX or Linux. If you use Windows, you might have to install a virtual machine to get a UNIX-like environment to continue with the rest of this instruction. A lot of this instruction is more verbose than needed to accomodate participants of different skill levels.

Please note that these are only optional. On the first day of this training, you will be provided with a link to a JupyterHub instance where the environment will be pre-made and ready to go!

2.1 0. Get Anaconda

Anaconda is a Python (and R) distribution that aims to provide everything needed for common scientific and machine learning situations out-of-the-box. We chose Anaconda for this tutorial as it significantly simplifies Python dependency management.

In practice, Anaconda can be used to manage different environment and packages. This setup document will assume that you have Anaconda installed as your default Python distribution.

You can download Anaconda here: <https://www.continuum.io/downloads>

After installing Anaconda, you can access its command-line interface with the `conda` command.

2.2 1. Create a new environment

Environments are a tool for sanitary software development. By this, we mean that you can install specific versions of packages without worrying that it breaks a dependency elsewhere.

Here is how you can create an environment with Anaconda

```
conda create -n dl4nlp python=3.6
```

2.3 2. Install Dependencies

2.3.1 2a. Activate the environment

After creating the environment, you need to **activate** the environment:

```
source activate dl4nlp
```

After an environment is activated, it might prepend/append itself to your console prompt to let you know it is active.

With the environment activated, any installation commands (whether it is `pip install X`, `python setup.py install` or using Anaconda's install command `conda install X`) will only install inside the environment.

2.3.2 2b. Install IPython and Jupyter

Two core dependencies are IPython and Jupyter. Let's install them first:

```
conda install ipython
conda install jupyter
```

To allow a jupyter notebooks to use this environment as their kernel, it needs to be linked:

```
python -m ipykernel install --user --name dl4nlp
```

2.3.3 2c. Installing CUDA (optional)

NOTE: CUDA is currently not supported out of the conda package control manager. Please refer to pytorch's github repository for compilation instructions.

If you have a CUDA compatible GPU, it is worthwhile to take advantage of it as it can significantly speedup training and make your PyTorch experimentation more enjoyable.

To install CUDA:

1. Download CUDA appropriate to your OS/Arch from [here](#).
2. Follow installation steps for your architecture/OS. For Ubuntu/x86_64, see [here](#).
3. Download and install CUDNN from [here](#).

Make sure you have the latest CUDA (8.0) and CUDNN (7.0).

2.3.4 2d. Install PyTorch

There are instructions on <http://pytorch.org> which detail how to install it. If you have been following along so far and have Anaconda installed with CUDA enabled, you can simply do:

```
conda install pytorch torchvision cuda80 -c soumith
```

The widget on PyTorch.org will let you select the right command line for your specific OS/Arch.

PLEASE NOTE. Make sure you have PyTorch 0.3.0. PyTorch has recently released version 0.4.0, but it has many code changes that we will not be incorporating at this time. The Anaconda installation method for this is:


```
conda install pytorch=0.3.1 torchvision -c pytorch
```

If you would like to install using pips and wheels:

```
pip install http://download.pytorch.org/whl/cpu/torch-0.3.1-cp36-cp36m-linux_x86_64.  
↪whl  
pip install torchvision
```

2.3.5 2e. Clone (or Download) Repository

At this point, you may have already cloned the tutorial repository. But if you have not, you will need it for the next step.

```
git clone https://github.com/joosthub/pytorch-nlp-tutorial-ny2018.git
```

If you do not have git or do not want to use it, you can also [download the repository as a zip file](#)

2.3.6 2f. Install Dependencies from Repository

Assuming the you have cloned (or downloaded and unzipped) the repository, please navigate to the directory in your terminal. Then, you can do the following:

```
pip install -r requirements.txt
```

Frequency Asked Questions

On this page, you will find a list of questions that we either anticipate people will ask or that we have been asked previously. They are intended to be the first stop for any confusion or trouble that might occur.

3.1 Do I Need to have a NVIDIA GPU enabled laptop?

Nope! While having a NVIDIA GPU enabled laptop will make the training run faster, we provide instructions for people who do not have one.

If you are plan on working on Natural Language Processing/Deep Learning in the future, a GPU enabled laptop might be a good investment.

CHAPTER 4

Migrating to PyTorch 0.4.0

The training session at O'Reilly AI in NYC, 2018 will be conducted using PyTorch 0.3.1. However, as of the end of April, PyTorch 0.4.0 has been released!

To help you understand how to migrate, the PyTorch folks have a wonderful migration guide found [here](#).

5.1 Problem 1

```
def f(x):
    if x.data[0] > 0:
        return torch.sin(x)
    else:
        return torch.cos(x)

x = torch.autograd.Variable(torch.FloatTensor([1]),
                             requires_grad=True)

y = f(x)
print(y)

y.backward()

x.grad

y.grad_fn
```

5.2 Problem 2

```
def cbow(phrase):
    words = phrase.split(" ")
    embeddings = []
    for word in words:
        if word in glove.word_to_index:
            embeddings.append(glove.get_embedding(word))
    embeddings = np.stack(embeddings)
    return np.mean(embeddings, axis=0)
```

```
cbow("the dog flew over the moon").shape

# >> (100,)

def cbow_sim(phrase1, phrase2):
    vec1 = cbow(phrase1)
    vec2 = cbow(phrase2)
    return np.dot(vec1, vec2) / (np.linalg.norm(vec1) * np.linalg.norm(vec2))

cbow_sim("green apple", "green apple")
# >> 1.0

cbow_sim("green apple", "apple green")
# >> 1.0

cbow_sim("green apple", "red potato")
# >> 0.749

cbow_sim("green apple", "green alien")
# >> 0.683

cbow_sim("green apple", "blue alien")
# >> 0.5799815958114477

cbow_sim("eat an apple", "ingest an apple")
# >> 0.9304712574359718
```

Warm Up Exercise

To get you back into the PyTorch groove, let's do some easy exercises. You will have 10 minutes. See how far you can get.

1. Use `torch.randn` to create two tensors of size (29, 30, 32) and (32, 100).
2. Use `torch.matmul` to matrix multiply the two tensors.
3. Use `torch.sum` on the resulting tensor, passing the optional argument of `dim=1` to sum across the 1st dimension. Before you run this, can you predict the size?
4. Create a new long tensor of size (3, 10) from the `np.random.randint` method.
5. Use this new long tensor to index into the tensor from step 3.
6. Use `torch.mean` to average across the last dimension in the tensor from step 5.

Fail Fast Prototype Mode

When building neural networks, you want things to either work or fail fast. Long iteration loops are the truest enemy of the machine learning practitioner.

To that end, the following techniques will help you out.

```
import torch
from torch import nn
from torch.autograd import Variable
# note, Variable deprecates in 0.4.0

# 2dim tensor.. aka a matrix
x = Variable(torch.randn(4, 5))

# this is the same as:
batch_size = 4
feature_size = 5
x = Variable(torch.randn(batch_size, feature_size))
```

You can construct whatever prototype variables you want doing this.

7.1 Prototyping an embedding

```
import torch
from torch import nn
from torch.autograd import Variable
# note, Variable deprecates in 0.4.0

batch_size = 4
sequence_size = 5
integer_range = 100
embedding_size = 25
# notice rand vs randn. rand is uniform (0,1), and randn is normal (-1,1)
random_numbers = torch.rand(batch_size, sequence_size) * integer_range
```

```
x = Variable(random_numbers.long())

embedder = nn.Embedding(num_embeddings=integer_range,
                        embedding_dim=embedding_size)

print(embedder(x).shape)
```

8.1 Exercise 1

```
import torch
from torch import nn
x = torch.randn(9, 10)
```

8.2 Exercise 2

```
import torch
from torch import nn

x2dim = torch.randn(9, 10)

# required and default parameters:
# fc = nn.Linear(in_features, out_features)
```

Task: Create a linear layer which works with x2dim

8.3 Exercise 3

```
import torch
from torch import nn

x3dim = torch.randn(9, 10, 11)

# required and default parameters:
# conv1 = nn.Conv1d(in_channels, out_channels, kernel_size, stride=1, padding=0)
```

Task: Create a convolution which works on x3dim

9.1 Exercise 1

```
indices = torch.arange(10).long()
indices = torch.from_numpy(np.random.randint(0, 10, size=(10,)))

emb = nn.Embedding(num_embeddings=100, embedding_dim=16)
emb(indices)
```

Task: Get the above code to work. Use the second indices method and change the size to a matrix (such as (10,11)).

9.2 Exercise 2

Task: Create a MultiEmbedding class which can input two sets of indices, embed them, and concat the results!

```
class MultiEmbedding(nn.Module):
    def __init__(self, num_embeddings1, num_embeddings2, embedding_dim1, embedding_
    ↪dim2):
        pass

    def forward(self, indices1, indices2):
        # use something like
        # z = torch.concat([x, y], dim=1)

        pass
```

Exercise: Interpolating Between Vectors

One fun option for the conditional generation code is to interpolate between the learned hidden vectors.

To do this, first look at the code for sampling given a specific nationality:

```

1 def sample_n_for_nationality(nationality, n=10, temp=0.8):
2     assert nationality in vectorizer.nationality_vocab.keys(), 'not a nationality we
↳trained on'
3     keys = [nationality] * n
4     init_vector = long_variable([vectorizer.nationality_vocab[key] for key in keys])
5     init_vector = net.conditional_emb(init_vector)
6     samples = decode_matrix(vectorizer,
7                             sample(net.emb, net.rnn, net.fc,
8                                   init_vector,
9                                   make_initial_x(n, vectorizer),
10                                  temp=temp))
11     return list(zip(keys, samples))

```

As you can see, we create a list of keys that is the length of the number of samples we want (n). And we use that list to retrieve the correct index from the vocabulary. Finally, we use that index in the conditional embedding inside the network to get the initial hidden state for the sampler.

To do this exercise, write a function that has the following signature:

```

def interpolate_n_samples_from_two_nationalities(nationality1, nationality2, weight,
↳n=10, temp=0.8):
    print('awesome stuff here')

```

This should retrieve the `init_vectors` for two different nationalities. Then, using the weight, combine the init vectors as $\text{weight} * \text{init_vector1} + (1 - \text{weight}) * \text{init_vector2}$.

For fun, after you finish this function, write a for loop which loops over the weight from 0.1 to 0.9 to see how it affects the generation.

Exercise: Sampling from an RNN

The goal of sampling from an RNN is to initialize the sequence in some way, feed it into the recurrent computation, and retrieve the next prediction.

To start, we create the initial vectors:

```
start_index = vectorizer.surname_vocab.start_index
batch_size = 2
# hidden_size = whatever hidden size the model is set to

initial_h = Variable(torch.ones(batch_size, hidden_size))
initial_x_index = Variable(torch.ones(batch_size).long()) * start_index
```

Then, we need to use these vectors to retrieve the next prediction:

```
# model is stored in variable called `net`

x_t = net.emb(initial_x_index)
print(x_t.shape)
h_t = net.rnn._compute_next_hidden(x_t, initial_h)

y_t = net.fc(h_t)
```

Now that we have a prediction vector, we can create a probability distribution and sample from it. Note we include a temperature hyper parameter for controlling how strongly we sample from the distribution (at high temperatures, everything is uniform, at low temperatures below 1, small differences are magnified). The temperature is always greater than 0.

```
temperature = 1.0
prediction_vector = F.softmax(y_t / temperature, dim=1)
x_index_t = torch.multinomial(y_t, 1)[: , 0]
```

Now we can start the cycle over again:

```
x_t = net.emb(x_index_t)
h_t = net.rnn._compute_next_hidden(x_t, h_t)
```

```
y_t = net.fc(h_t)
```

Write a for loop which repeats this sequence and appends the `x_t` variable to a list.

Then, we can do the following:

```
final_x_indices = torch.stack(x_indices).squeeze().permute(1, 0)

# stop here if you don't know what cpu, data, and numpy do. Ask away!
final_x_indices = final_x_indices.cpu().data.numpy()

# loop over the items in the batch
results = []
for i in range(len(final_x_indices)):
    tokens = []
    index_vector = final_x_indices[i]
    for x_index in index_vector:
        if vectorizer.surname_vocab.start_index == x_index:
            continue
        elif vectorizer.surname_vocab.end_index == x_index:
            break
        else:
            token = vectorizer.surname_vocab.lookup(x_index)
            tokens.append(token)

    sampled_surname = "".join(tokens)
    results.append(sampled_surname)
    tokens = []
```

Design Pattern: Attention

Attention is a useful pattern for when you want to take a collection of vectors—whether it be a sequence of vectors representing a sequence of words, or an unordered collections of vectors representing a collection of attributes—and summarize them into a single vector. This has similar analogs to the CBOW examples we saw on Day 1, but instead of just averaging or using max pooling, we are learning a function which learns to compute the weights for each of the vectors before summing them together.

Importantly, the weights that the attention module is learning is a valid probability distribution. This means that weighting the vectors by the value the attention module learns can additionally be seen as computing the Expectation. Or, it could as interpolating. In any case, attention’s main use is to select ‘softly’ amongst a set of vectors.

The attention vector has several different published forms. The one below is very simple and just learns a single vector as the attention mechanism.

Using the `new_parameter` function we have been using for the RNN notebooks:

```
def new_parameter(*size):
    out = Parameter(FloatTensor(*size))
    torch.nn.init.xavier_normal(out)
    return out
```

We can then do:

```
class Attention(nn.Module):
    def __init__(self, attention_size):
        super(Attention, self).__init__()
        self.attention = new_parameter(attention_size, 1)

    def forward(self, x_in):
        # after this, we have (batch, dim1) with a diff weight per each cell
        attention_score = torch.matmul(x_in, self.attention).squeeze()
        attention_score = F.softmax(attention_score).view(x_in.size(0), x_in.size(1), ↵
↵1)
        scored_x = x_in * attention_score

        # now, sum across dim 1 to get the expected feature vector
```

```
        condensed_x = torch.sum(scored_x, dim=1)

        return condensed_x

attn = Attention(100)
x = Variable(torch.randn(16, 30, 100))
attn(x).size() == (16, 100)
```

For participants of the Training Tutorial in NY, please fill out this form! <https://goo.gl/forms/iLRlpoutWBy3As8Q2>

Hello! This is a directory of resources for a training tutorial to be given at the O'Reilly AI Conference in New York City on Sunday, April 29, and Monday, April 30, 2018.

Please read below for general information. You can find the github repository at [this link](#). Please note that there are two ways to engage in this training (described below).

More information will be added to this site as the training progresses. Specifically, we will be adding a 'recipes' section, 'errata' section, and a 'bonus exercise' section as the training progresses!

13.1 Prerequisites:

- A working knowledge of Python and the command line
- Familiarity with precalc math (multiply matrices, dot products of vectors, etc.) and derivatives of simple functions (If you are new to linear algebra, this video course is handy.)
- A general understanding of machine learning (setting up experiments, evaluation, etc.) (useful but not required)

13.2 Hardware and/or installation requirements:

- **There are two options:**
 1. **Using O'Reilly's online resources.** For this, you only need a laptop; on the first day, we will provide you with credentials and a URL to use an online computing resource (a JupyterHub instance) provided by O'Reilly. You will be able to access Jupyter notebooks through this and they will persist until the end of the second day of training (April 30th). This option is not limited by what operating system you have. You will need to have a browser installed.
 2. **Setting everything up locally.** For this, you need a laptop with the PyTorch environment set up. This is only recommended if you want to have the environment locally or have a laptop with a GPU. (If you have trouble following the provided instructions or if you find any mistakes, please file an issue [here](#).) This option is limited to Macs and Linux users only (sorry Windows users!). Be sure you check the [LOCAL_RUN_README.md](#).