



Sequence Tagging with LSTMs

CS 780/880 Natural Language Processing Lecture 16

Samuel Carton, University of New Hampshire

Last lecture

RNNs


- One-to-one
- **Many-to-one**
- Many-to-many

LSTMS

Increasing RNN capacity

- Depth
- Bidirectionality


Dropout



LSTMs: NLP Swiss army knife

LSTMs are exciting for us because they are the Swiss army knife of NLP models.


- Sequence classification
- Sequence tagging
- Language modeling
- Text-to-text (e.g. translation)



LSTMs: NLP Swiss army knife

LSTMs are exciting for us because they are the Swiss army knife of NLP models.


- **Sequence classification**
- Sequence tagging
- Language modeling
- Text-to-text (e.g. translation)



LSTMs: NLP Swiss army knife

LSTMs are exciting for us because they are the Swiss army knife of NLP models.


- Sequence classification
- **Sequence tagging**
- Language modeling
- Text-to-text (e.g. translation)



LSTMs: NLP Swiss army knife

LSTMs are exciting for us because they are the Swiss army knife of NLP models.


- Sequence classification
- Sequence tagging
- **Language modeling**
- Text-to-text (e.g. translation)



LSTMs: NLP Swiss army knife

LSTMs are exciting for us because they are the Swiss army knife of NLP models.


- Sequence classification
- Sequence tagging
- Language modeling
- **Text-to-text (e.g. translation)**



LSTMs: NLP Swiss army knife

LSTMs are exciting for us because they are the Swiss army knife of NLP models.

- Sequence classification
- **Sequence tagging**
- Language modeling
- Text-to-text (e.g. translation)



Sequence tagging

Basic idea: Given a corpus of text where each **word** has a label, learn to predict word labels for unseen texts

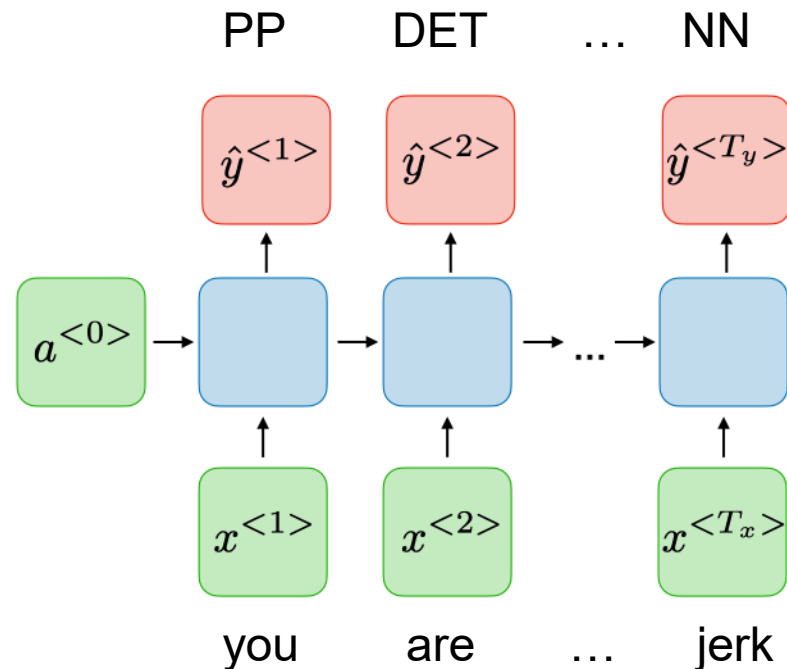
- Part-of-speech tagging
- Named entity recognition
 - “In his speech to the UN today, **George Bush** addressed the rising problems of...”
- Explanations
 - “You are a real **piece of garbage** human being.” → Predicted toxic



Sequence tagging

Context sensitive.

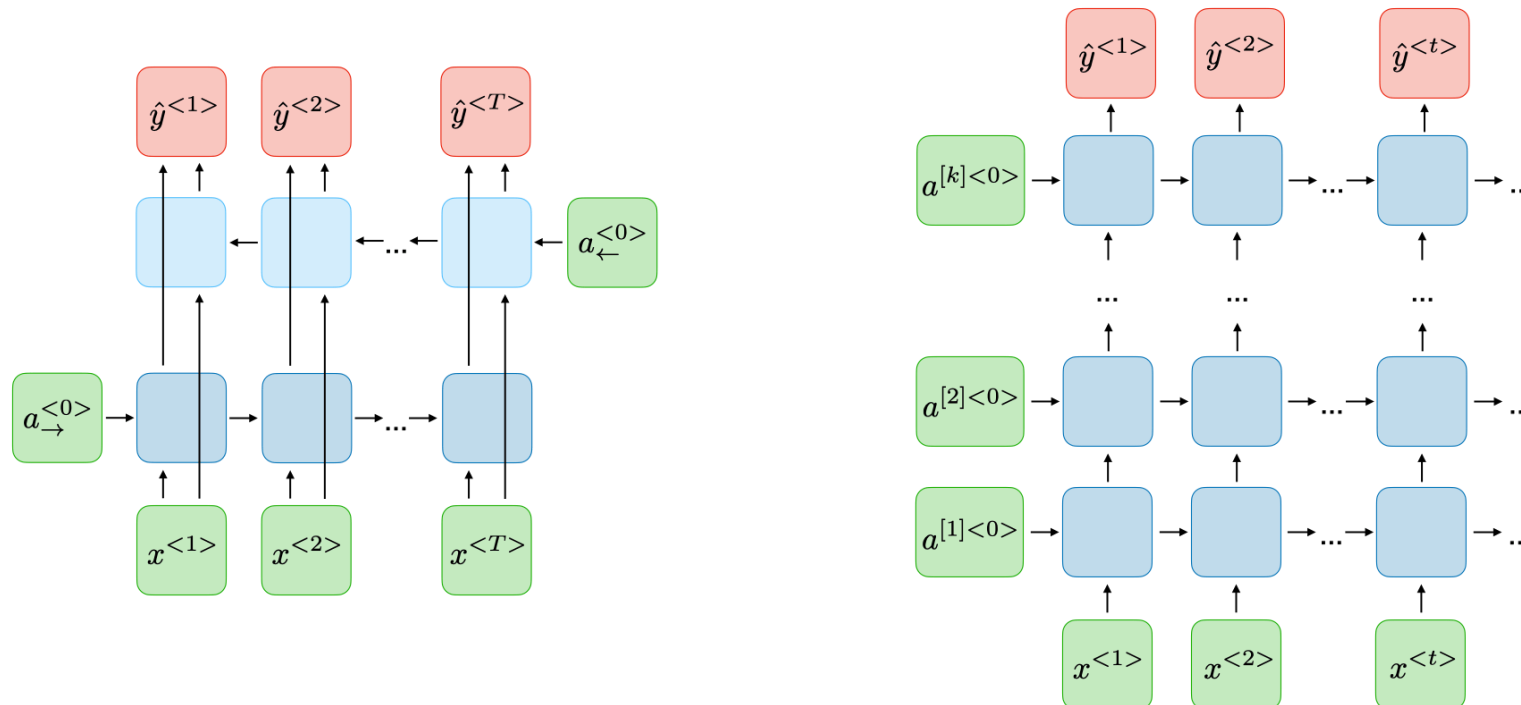
- “You are a real jerk!”
- “I am really craving some Jamaican jerk chicken right now.”



Sequence tagging

Context sensitive.

- “You are a real jerk!”
- “I am really craving some Jamaican jerk chicken right now.”



Loading GloVe vectors with Gensim

```
1 import gensim.downloader as api

1 # We're gonna be working with Twitter text today, so we'll use a Twitter-specific set of
2 # pretrained word embeddings
3 vector_model = api.load('glove-twitter-100')

[=====] 99.9% 386.6/387.1MB downloaded
```



Loading GloVe vectors with Gensim

```
1 # We can see that this particular model has special embeddings for
2 # various kinds of things you'll find in tweets
3 vector_model.index_to_key[0:15]
```

```
['<user>',
 '.',
 ':',
 'rt',
 ',',
 '<repeat>',
 '<hashtag>',
 '<number>',
 '<url>',
 '!',
 'i',
 'a',
 '"',
 'the',
 '?']
```

```
1 # It's also apparently got multilingual stuff in it
2 vector_model.index_to_key[-15:]
```

```
['game',
 'アマネシア',
 'II',
 'カリカリ',
 'キイ',
 'ゲッ',
 'テハロ',
 'テモ',
 'ハインイ',
 'ハンチ',
 'ヤマタマ',
 'ヨイヨ',
 'オヤスミ',
 '<unk>',
 '<pad>']
```



Twitter POS tagging dataset

Named Entity Recognition in Tweets: An Experimental Study (Ritter et al., 2011)

https://raw.githubusercontent.com/aritter/twitter_nlp/master/data/annotated/pos.txt

```
@paulwalk USR  
It PRP  
's VBZ  
the DT  
view NN  
from IN  
where WRB  
I PRP  
'm VBP  
living VBG  
for IN  
two CD  
weeks NNS  
. .  
Empire NNP  
State NNP  
Building NNP  
= SYM  
ESB NNP  
. .  
Pretty RB  
bad JJ  
storm NN  
here RB  
last JJ  
evening NN  
. .
```

```
@Miss_SOTO USR  
I PRP  
think VBP  
that DT  
's VBZ  
when WRB  
I PRP  
'm VBP  
gonna VBG  
be VB  
there RB  
  
On IN  
Thanksgiving NNP  
after IN  
you PRP  
done VBN  
eating VBG  
its PRP  
#TimeToGetOut HT  
unless IN  
you PRP  
wanna VBP  
help VB  
with IN  
the DT  
dishes NNS
```



Reading and preprocessing POS data

```
@paulwalk USR  
It PRP  
's VBZ  
the DT  
view NN  
from IN  
where WRB  
I PRP  
'm VBP  
living VBG  
for IN  
two CD  
weeks NNS  
. .  
Empire NNP  
State NNP  
Building NNP  
= SYM  
ESB NNP  
. .  
Pretty RB  
bad JJ  
storm NN  
here RB  
last JJ  
evening NN  
. .
```

```
1 import pandas as pd  
2 import csv  
3 raw_pos_df = pd.read_csv(pos_url, sep=' ', quoting=csv.QUOTE_NONE, names=['token', 'tag'])  
4 display(raw_pos_df)
```

	token	tag
0	@paulwalk	USR
1	It	PRP
2	's	VBZ
3	the	DT
4	view	NN
...
15180	wanna	VBP
15181	talk	VB
15182	to	TO
15183	u	PRP
15184	!!!!	.

15185 rows x 2 columns



Reading and preprocessing POS data

```
6 tagged_tweets = []
7 current_tweet = {'tokens':[], 'tags':[]}
8
9 for row_index, row in raw_pos_df.iterrows(): # this will yield each row as a pd.Series object
10     if row['token'].startswith('@'): # if we hit a new tweet...
11         if len(current_tweet['tokens']) > 0: # add the current tweet to the list if it isn't empty
12             tagged_tweets.append(current_tweet)
13             current_tweet = {'tokens':[], 'tags':[]} #then reset the current tweet
14
15     current_tweet['tokens'].append(row['token']) # then begin accumulating into current tweet
16     current_tweet['tags'].append(row['tag'])
17
18 if len(current_tweet['tokens']) > 0: # add the current tweet to the list if it isn't empty
19     tagged_tweets.append(current_tweet)
20
21 #Pandas knows how to create a DataFrame from a list of dictionaries
22 pos_df = pd.DataFrame(tagged_tweets)
23 display(pos_df)
```



Reading and preprocessing POS data

	tokens	tags
0	[@paulwalk, It, 's, the, view, from, where, I,...	[USR, PRP, VBZ, DT, NN, IN, WRB, PRP, VBP, VBG...
1	[@MISS_SOTO, I, think, that, 's, when, I, 'm, ...	[USR, PRP, VBP, DT, VBZ, WRB, PRP, VBP, VBG, V...
2	[@robmoyses, Eyeopener, vs, ., Ryerson, Quiddi...	[USR, NNP, CC, ., NNP, NN, NN, DT, NNP, IN, CD...
3	[@RUQuidditch, #Rams, RT]	[USR, HT, RT]
4	[@ZodiacFacts, :, #ZodiacFacts, As, an, #Aries...	[USR, :, HT, IN, DT, HT, NN, VBZ, RB, IN, WP, ...
...
467	[@DailyCaller, tomorrow, !, http://is.gd/fKm4j...	[USR, NN, ., URL, PRP, VBP, RB, VBN, TO, NNP, ...
468	[@DORSEY33, lol, aw, ., i, thought, u, was, ta...	[USR, UH, UH, ., PRP, VBD, PRP, VBD, VBG, IN, ...
469	[@Bibhu2109, No, ., He, 's, gonna, run, out, o...	[USR, UH, ., PRP, VBZ, VBG, VB, IN, IN, NN, DT...
470	[@SincerelyKRenee, but, u, can, just, get, her...	[USR, CC, PRP, MD, RB, VB, PRP, NN, NN, CC, VB...
471	[@MyssLidia, :, If, u, call, someone, 5x, 's, ...	[USR, :, IN, PRP, VBP, NN, NN, VBZ, DT, NN, CC...

472 rows x 2 columns



Reading and preprocessing POS data

49 possible POS tags in this particular dataset

```
1 # And we can also get a canonical list of all possible tags
2 tags = raw_pos_df['tag'].unique()
3 tags

array(['USR', 'PRP', 'VBZ', 'DT', 'NN', 'IN', 'WRB', 'VBP', 'VBG', 'CD',
      'NNS', '.', 'NNP', 'SYM', 'RB', 'JJ', ':', 'URL', 'HT', 'VB',
      'VBN', 'RT', 'CC', 'TO', 'WP', ',', 'UH', 'RP', 'JJS', 'PRP$',
      'VBD', "'", 'POS', 'JJR', 'MD', 'NNPS', '(', 'WDT', 'VPP', 'EX',
      ')', 'PDT', 'RBR', 'LS', 'TD', 'FW', 'RBS', 'NONE', 'O'],
      dtype=object)
```



Reading and preprocessing POS data

```
3 def token_to_ID(token):
4     token = token.lower()
5     if token.startswith('@'):
6         return vector_model.key_to_index['<user>']
7     elif token.startswith('#'):
8         return vector_model.key_to_index['<hashtag>']
9     elif token.startswith('http'):
10        return vector_model.key_to_index['<url>']
11    elif token in vector_model.key_to_index:
12        return vector_model.key_to_index[token]
13    else:
14        return vector_model.key_to_index['<unk>']
15
16 pos_df['input_ids'] = pos_df['tokens'].apply(lambda tokens:[token_to_ID(token) for token in tokens])
```

```
1 # We also need to map tags to tag IDs
2 tag2id = {tag:id for id, tag in enumerate(tags)}
3 pos_df['tag_ids'] = pos_df['tags'].apply(lambda tags:[tag2id[tag] for tag in tags])
```



Reading and preprocessing POS data

	tokens	tags	input_ids	tag_ids
0	[@paulwalk, It, 's, the, view, from, where, I,...	[USR, PRP, VBZ, DT, NN, IN, WRB, PRP, VBP, VBG...	[0, 33, 41, 13, 3056, 133, 329, 10, 57, 1704, ...	[0, 1, 2, 3, 4, 5, 6, 1, 7, 8, 5, 9, 10, 11, 1...
1	[@MISS_SOTO, I, think, that, 's, when, I, 'm, ...	[USR, PRP, VBP, DT, VBZ, WRB, PRP, VBP, VBG, V...	[0, 10, 186, 45, 41, 92, 10, 57, 316, 56, 175,...	[0, 1, 7, 3, 2, 6, 1, 7, 8, 19, 14, 5, 12, 5, ...
2	[@robmoyses, Eyeopener, vs, ., Ryerson, Quiddi...	[USR, NNP, CC, ., NNP, NN, NN, DT, NNP, IN, CD...	[0, 519575, 917, 1, 215106, 85242, 302, 53, 12...	[0, 12, 22, 11, 12, 4, 4, 3, 12, 5, 9, 4, 4, 1...
3	[@RUQuidditch, #Rams, RT]	[USR, HT, RT]	[0, 6, 5]	[0, 18, 21]
4	[@ZodiacFacts, ., #ZodiacFacts, As, an, #Aries...	[USR, ., HT, IN, DT, HT, NN, VBZ, RB, IN, WP, ...	[0, 2, 6, 124, 172, 6, 6315, 32, 44, 121, 86, ...	[0, 16, 18, 5, 3, 18, 4, 2, 14, 5, 24, 1, 7, 5...
...
467	[@DailyCaller, tomorrow, I, http://is.gd/fKm4j...	[USR, NN, ., URL, PRP, VBP, RB, VBN, TO, NNP, ...	[0, 328, 9, 8, 10, 57, 55, 4968, 16, 218, 110,...	[0, 4, 11, 17, 1, 7, 14, 20, 23, 12, 14, 5, 1,...
468	[@DORSEY33, lol, aw, ., i, thought, u, was, ta...	[USR, UH, UH, ., PRP, VBD, PRP, VBD, VBG, IN, ...	[0, 88, 751, 1, 10, 621, 51, 93, 3427, 734, 59...	[0, 26, 26, 11, 1, 30, 1, 30, 8, 5, 3, 4, 11, ...
469	[@Bibhu2109, No, ., He, 's, gonna, run, out, o...	[USR, UH, ., PRP, VBZ, VBG, VB, IN, IN, NN, DT...	[0, 30, 1, 107, 41, 316, 899, 99, 39, 580, 191...	[0, 26, 11, 1, 2, 8, 19, 5, 5, 4, 3, 4, 26]
470	[@SincerelyKRenee, but, u, can, just, get, her...	[USR, CC, PRP, MD, RB, VB, PRP, NN, NN, CC, VB...	[0, 79, 51, 102, 59, 87, 168, 185, 148, 36, 43...	[0, 22, 1, 34, 14, 19, 1, 4, 4, 22, 7, 14, 19,...
471	[@MyssLidia, ., If, u, call, someone, 5x, 's, ...	[USR, ., IN, PRP, VBP, NN, NN, VBZ, DT, NN, CC...	[0, 2, 74, 51, 462, 238, 1193514, 41, 11, 125,...	[0, 16, 5, 1, 7, 4, 4, 2, 3, 4, 22, 1, 7, 14, ...

472 rows x 4 columns



Training a LSTM POS tagger—Dataset

```
6 class POSTaggingDataset(Dataset):
7     def __init__(self,
8                 tag_ids=None,
9                 input_ids=None):
10
11         self.tag_ids = tag_ids
12         self.input_ids = input_ids
13
14     def __len__(self):
15         return len(self.tag_ids)
16
17     def __getitem__(self, idx):
18         rdict = {
19             'tag_ids': torch.tensor(self.tag_ids[idx], dtype=torch.int64),
20             'input_ids': torch.tensor(self.input_ids[idx], dtype=torch.int64)
21         }
22         return rdict

```

```
1 train_dataset = POSTaggingDataset(train_df['tag_ids'], train_df['input_ids'])
2 dev_dataset = POSTaggingDataset(dev_df['tag_ids'], dev_df['input_ids'])

```



Training a LSTM POS tagger—Dataset

```
1 from pprint import pprint
2 pprint(train_dataset[0])
3 print(train_dataset[0]['input_ids'].shape)

{'input_ids': tensor([
    0, 2, 525, 291, 99, 28, 80, 140,
  10510, 46, 3761, 53, 435, 9, 183, 538,
   16, 12446, 55, 1898, 1417, 15, 68, 291,
   8, 59, 121, 231, 320, 9, 211, 16,
  960077, 389, 1193514, 1352, 1, 2431, 80, 143,
   6404, 41, 3645, 35, 13, 4948, 148, 1]),
 'tag_ids': tensor([ 0, 16, 19, 1, 27, 16, 1, 7, 20, 5, 12, 3, 4, 11, 19, 15, 23, 19,
   5, 12, 2, 1, 7, 1, 17, 14, 14, 14, 4, 11, 8, 23, 12, 14, 26, 26,
  11, 14, 1, 30, 12, 32, 4, 5, 3, 4, 4, 11])}
torch.Size([48])
```



Training a LSTM POS tagger—DataLoader

```
3 first_train_batch = next(iter(train_dataloader))
4 print('First training batch:')
5 pprint(first_train_batch)
6
7 print('First training batch sizes:')
8 pprint({key:value.shape for key, value in first_train_batch.items()})
```

First training batch:

```
{'input_ids': tensor([[ 0, 10, 247, ..., 1193515, 1193515, 1193515],
 [ 0, 10, 64, ..., 1193514, 8, 5],
 [ 0, 277, 6, ..., 1193515, 1193515, 1193515],
 ...,
 [ 0, 122, 524, ..., 1193515, 1193515, 1193515],
 [ 0, 265, 21, ..., 1193515, 1193515, 1193515],
 [ 0, 2, 6, ..., 1193515, 1193515, 1193515]]),
 'tag_ids': tensor([[ 0, 1, 7, ..., 0, 0, 0],
 [ 0, 1, 7, ..., 16, 17, 21],
 [ 0, 29, 18, ..., 0, 0, 0],
 ...,
 [ 0, 15, 4, ..., 0, 0, 0],
 [ 0, 19, 1, ..., 0, 0, 0],
 [ 0, 16, 18, ..., 0, 0, 0]])}
```

First training batch sizes:

```
{'input_ids': torch.Size([10, 66]), 'tag_ids': torch.Size([10, 66])}
```



Training a LSTM POS tagger—Model

```
1 ! pip install --quiet "pytorch-lightning==1.9.4"
2
3 # PyTorch Lightning recently released v2.0 (March 15 2023), but it changes some syntax,
4 # so I am teaching the last 1.9.x version for now.
5 # https://github.com/Lightning-AI/lightning/releases
```

827.8/827.8 KB	39.2 MB/s	eta 0:00:00
519.2/519.2 KB	41.2 MB/s	eta 0:00:00
1.0/1.0 MB	62.7 MB/s	eta 0:00:00
264.6/264.6 KB	28.7 MB/s	eta 0:00:00
114.2/114.2 KB	14.6 MB/s	eta 0:00:00
158.8/158.8 KB	18.6 MB/s	eta 0:00:00



Training a LSTM POS tagger—Model

```
6 class LSTMPOSTagger(pl.LightningModule):
7     def __init__(self,
8                 word_vectors:np.ndarray,
9                 num_classes:int,
10                learning_rate:float,
11                padding_id:int,
12                lstm_hidden_size:int=100, # how big the inner vectors of the LSTM will be,
13                lstm_layers:int =2, # how many layers the LSTM will have
14                dropout_prob:float=0.1,
15                **kwargs):
16         super().__init__( **kwargs)
17
18         # The __init__ function will be identical to the classifier version
19         self.word_embeddings = torch.nn.Embedding.from_pretrained(torch.tensor(word_vectors),
20                                                                freeze=True)
21         self.lstm = torch.nn.LSTM(input_size = word_vectors.shape[1], # The LSTM will be taking in word vectors
22                                 hidden_size = lstm_hidden_size,
23                                 num_layers=lstm_layers,
24                                 bidirectional=True,
25                                 dropout=dropout_prob,
26                                 batch_first=True # This is important. Set to False by default for some reason.
27                                 )
28
29         # Output layer input size has to be doubled because the LSTM is bidirectional
30         self.output_layer = torch.nn.Linear(2*lstm_hidden_size, num_classes)
31         self.lstm_layers = lstm_layers
32         self.learning_rate = learning_rate
33         self.padding_id = padding_id # we'll need this later
34         self.train_accuracy = Accuracy(task='multiclass', num_classes=num_classes)
35         self.val_accuracy = Accuracy(task='multiclass', num_classes=num_classes)
```



Training a LSTM POS tagger—Model

```
def forward(self, tag_ids:torch.Tensor, input_ids:torch.Tensor, verbose=False):  
  
    #The first part of the forward() function is the same too  
    inputs_embeds = self.word_embeddings(input_ids) #(batch size x sequence length x embedding size)  
    padding_mask = (input_ids != self.padding_id).int()  
    input_lengths = padding_mask.sum(dim=1).detach().cpu()  
    packed_embeddings = pack_padded_sequence(inputs_embeds, input_lengths, batch_first=True, enforce_sorted=False)  
    packed_output, (final_hidden, final_state) = self.lstm.forward(packed_embeddings)  
  
    # But now we need to look at all the LSTM output, not just the final hidden state  
    # So first we unpack the packed output  
    output, _ = pad_packed_sequence(packed_output, batch_first=True, padding_value=0.0, total_length=input_ids.shape[1])  
  
    # output is actually nicely shaped for us: (batch size x sequence length x 2*lstm hidden size)  
    py_logits = self.output_layer(output) #(batch size x sequence length x num_classes)  
    py = torch.argmax(py_logits, dim=2)  
  
    # We end up with one loss value per token  
    # Annoyingly, this function wants the class to be the second dimension  
    losses = torch.nn.functional.cross_entropy(py_logits.transpose(1,2), tag_ids, reduction='none')  
  
    # Then the final thing we need to do is zero out the losses for padding  
    padded_losses = losses * padding_mask  
    loss = padded_losses.mean()  
  
    return {'py':py,  
            'loss':loss}
```



Training a LSTM POS tagger—Model

```
65 # And then everything else is the same!
66 def configure_optimizers(self):
67     return [torch.optim.Adam(self.parameters()), lr=self.learning_rate]
68
69 def training_step(self, batch, batch_idx):
70     result = self.forward(**batch)
71     loss = result['loss']
72     self.log('train_loss', result['loss'])
73     self.train_accuracy.update(result['py'], batch['tag_ids'])
74     return loss
75
76 def training_epoch_end(self, outs):
77     print(f'Epoch {self.current_epoch} training accuracy:', self.train_accuracy.compute())
78     self.train_accuracy.reset()
79
80 def validation_step(self, batch, batch_idx):
81     result = self.forward(**batch)
82     self.val_accuracy.update(result['py'], batch['tag_ids'])
83     return result['loss']
84
85 def validation_epoch_end(self, outs):
86     print(f'Epoch {self.current_epoch} validation accuracy:', self.val_accuracy.compute())
87     self.val_accuracy.reset()
```



Training a LSTM POS tagger—Model

```
1 tagger_model = LSTMPOSTagger(word_vectors=vector_model.vectors,
2                               num_classes = len(tags),
3                               learning_rate = 0.001, #I'll typically start with something like 1e-3 for LSTMs
4                               padding_id = vector_model.key_to_index['<pad>'],
5                               lstm_hidden_size=100,
6                               lstm_layers=2,
7                               dropout_prob=0.1)
8 print('Model:')
9 print(tagger_model)
```

Model:

```
LSTMPOSTagger(
  (word_embeddings): Embedding(1193516, 100)
  (lstm): LSTM(100, 100, num_layers=2, batch_first=True, dropout=0.1, bidirectional=True)
  (output_layer): Linear(in_features=200, out_features=49, bias=True)
  (train_accuracy): MulticlassAccuracy()
  (val_accuracy): MulticlassAccuracy()
)
```



Training a LSTM POS tagger—Model

```
1 from pprint import pprint
2 with torch.no_grad():
3     first_train_output = tagger_model(**first_train_batch, verbose=True)
4
5 print('First training output:')
6 pprint(first_train_output)
7
8 print('Output item shapes:')
9 pprint({key:value.shape for key, value in first_train_output.items()})
```

```
First training output:
{'loss': tensor(0.2590),
 'py': tensor([[ 0,  1,  7, ..., 12, 12, 12],
               [ 0,  1,  7, ..., 16, 17, 21],
               [ 0, 29, 18, ..., 12, 12, 12],
               ...,
               [ 0, 12,  4, ..., 12, 12, 12],
               [ 0, 19,  1, ..., 12, 12, 12],
               [ 0, 16, 18, ..., 12, 12, 12]])}
Output item shapes:
{'loss': torch.Size([]), 'py': torch.Size([10, 66])}
```



Training a LSTM POS tagger—Trainer

```
1 from pytorch_lightning import Trainer
2 from pytorch_lightning.callbacks.progress import TQDMProgressBar
3
4 pos_trainer = Trainer(
5     accelerator="auto",
6     devices=1 if torch.cuda.is_available() else None,
7     max_epochs=10,
8     callbacks=[TQDMProgressBar(refresh_rate=20)],
9     val_check_interval = 0.5,
10 )
```

```
INFO:pytorch_lightning.utilities.rank_zero:GPU available: True (cuda), used: True
INFO:pytorch_lightning.utilities.rank_zero:TPU available: False, using: 0 TPU cores
INFO:pytorch_lightning.utilities.rank_zero:IPU available: False, using: 0 IPUs
INFO:pytorch_lightning.utilities.rank_zero:HPU available: False, using: 0 HPUs
```



Training a LSTM POS tagger—Trainer

```
1 pos_trainer.fit(model=tagger_model,  
2 | | | | | train_data loaders=train_data loader,  
3 | | | | | val_data loaders=dev_data loader)
```

```
Epoch 0 validation accuracy: tensor(0.0449, device='cuda:0')  
Epoch 0 validation accuracy: tensor(0.0723, device='cuda:0')  
Epoch 0 training accuracy: tensor(0.0509, device='cuda:0')  
Epoch 1 validation accuracy: tensor(0.1089, device='cuda:0')  
Epoch 1 validation accuracy: tensor(0.1472, device='cuda:0')  
Epoch 1 training accuracy: tensor(0.1122, device='cuda:0')  
Epoch 2 validation accuracy: tensor(0.1831, device='cuda:0')  
Epoch 2 validation accuracy: tensor(0.2083, device='cuda:0')  
Epoch 2 training accuracy: tensor(0.1895, device='cuda:0')  
Epoch 3 validation accuracy: tensor(0.2194, device='cuda:0')  
Epoch 3 validation accuracy: tensor(0.2348, device='cuda:0')  
Epoch 3 training accuracy: tensor(0.2285, device='cuda:0')  
Epoch 4 validation accuracy: tensor(0.2426, device='cuda:0')  
Epoch 4 validation accuracy: tensor(0.2483, device='cuda:0')  
Epoch 4 training accuracy: tensor(0.2610, device='cuda:0')  
Epoch 5 validation accuracy: tensor(0.2530, device='cuda:0')  
Epoch 5 validation accuracy: tensor(0.2626, device='cuda:0')  
Epoch 5 training accuracy: tensor(0.2768, device='cuda:0')  
Epoch 6 validation accuracy: tensor(0.2643, device='cuda:0')  
Epoch 6 validation accuracy: tensor(0.2719, device='cuda:0')  
Epoch 6 training accuracy: tensor(0.2812, device='cuda:0')  
Epoch 7 validation accuracy: tensor(0.2737, device='cuda:0')  
Epoch 7 validation accuracy: tensor(0.2753, device='cuda:0')  
Epoch 7 training accuracy: tensor(0.3019, device='cuda:0')  
Epoch 8 validation accuracy: tensor(0.2769, device='cuda:0')  
Epoch 8 validation accuracy: tensor(0.2758, device='cuda:0')  
Epoch 8 training accuracy: tensor(0.3055, device='cuda:0')  
Epoch 9 validation accuracy: tensor(0.2744, device='cuda:0')  
Epoch 9 validation accuracy: tensor(0.2781, device='cuda:0')  
INFO:pytorch_lightning.utilities.rank_zero:`Trainer.fit` stopped: `max_epochs=10` reached.  
Epoch 9 training accuracy: tensor(0.3073, device='cuda:0')
```


Concluding thoughts

Sequence tagging

- POS tagging

LSTMs as a NLP Swiss army knife

Domain-specific word embeddings

Masked loss

