

Assignment 3 - LSTMs and Image Synthesis

Due **April 25, 2018**

1 LSTM Language Generation

Start with the code from the last assignment but this time around, use an LSTM model. Do you get better results than with the CNN? If so, think of a reason why. If not, think of a reason why not. **Submit your code and a sample of generated output.**

2 Seq2Seq Image Synthesis

We will take a look at a simple image synthesis model. You will implement a seq2seq model that takes as input a digit (i.e. values from 0 through 9) and returns an image that resembles a hand-drawn digit (the inverse of the mnist problem we've seen in class so far).

Today's state of the art approaches for this problem don't use seq2seq models but it helps to have a diverse arsenal at your disposal. We have some starter code available at http://deeplearningcourse.onai.com/files/sample_code.py. This starter code reads the training data and at the end of each epoch saves a generated images (one per digit) on disk.

You are allowed to use any framework (Tensorflow, PyTorch, Keras, etc). The provided solution will be in Keras.

The following questions will help you reason and piece together your model.

- What does your input look like? (Note this is a digit but you can represent a digit in a variety of ways so think about the shape of your input.)
- What does your output look like? (Note that the output is an image with the dimensions of an image in the MNIST dataset.)
- How quickly did your solution converge?
- Does your output look reasonable (in your judgment)?

Submit your code and an example or two of generated images.

3 Administrative

You are allowed to work in groups of up to 3 in size. Turn in a hard copy of your submission in class on April 25.

Also email a pdf of this submission (and your Homework 2 submission if you have not already done so) to guha.jayachandran@sjsu.edu. Strong submissions will be posted on the website unless you object.