


# Homework assignments for Class #2

1. You are training a neural network using the gradient descent. You chose a particular learning rate and found that the training loss decreases as more epochs were performed (the blue curve). You are not sure if this is the best you could do. So, you tried a *smaller* learning rate. Which of the following red loss curves might you observe now?



2. The first layer of the DanQ model takes as input one-hot encoded DNA fragment of length 1000bp, and convolves it with 320 one-dimensional filters of size 26, using “valid” padding and  $\text{strides}=1$ . What is the shape of the output from this layer?

3. Implement on Biowulf a simple program that makes use of the LSTM layer. To this end, copy the script `simplernn_train.py` to your local folder and use it as a code base. Replace “SimpleRNN” for “LSTM” everywhere in the code. Add the “summary” command as described in the lecture. Rename the modified script to “`lstm_train.py`”. Make sure the script runs on Biowulf (command: `./lstm_train.py`). How many parameters have been used by the LSTM layer with 7 channels?