

# Ling 575j hw6

Due 11PM on May 11, 2023

In this assignment, you will

- Develop understanding of recurrent neural networks
- Implement components of data processing
- Implement key pieces of two variants of a recurrent model architecture

All files referenced herein may be found in `/dropbox/22-23/575j/hw6/` on patas.

## 1 Recurrent Neural Network Encoders [30 pts]

### Q1: Understanding RNNs [10 pts]

- What is the main limitation of feed-forward neural networks that is overcome by recurrent networks, and how do recurrent networks achieve this? [6 pts]
- The Vanilla RNN equation has the form  $h_t = f(h_{t-1}, x_t)$ . What extra ‘ingredient’ does the LSTM add to this general form? What problem is the LSTM designed to solve? [4 pts]

**Q2: LSTM Update** One of the “central” equations in the LSTM computation is the following:

$$c_t = f_t \odot c_{t-1} + i_t \odot \hat{c}_t$$

This equation performs an essential update of one part of the LSTM. Please answer: [10 pts]

- What is  $c_t$ ?
- What is the range of  $f_t$  and what is its purpose?
- What is the range of  $i_t$  and what is its purpose?
- What is  $\hat{c}_t$ ?
- In your own words, describe how this equation implements the central “update” inside of an LSTM.

**Q3: Counting parameters** Let  $d_e$  be the dimension of word embeddings and  $d_h$  the hidden state size. Focusing on just the recurrent cell (and so ignoring the embedding and output layers): [10 pts]

- How many parameters are there in a Vanilla RNN cell? [4 pts]
- How many parameters are there in an LSTM cell? [6 pts]

Note: for this problem, you can assume that the RNN cell is at the ‘bottom’ of a possibly-deep RNN, so the inputs to the cell are word embeddings, not earlier layers’ hidden states.

## 2 Implementing an RNN Sentiment Classifier [30 pts]

In the coding portion of this assignment, you will implement (components of) a classifier for the Stanford Sentiment Treebank, using RNNs as encoders. In particular, the model will take the final hidden state of an RNN that has read reviews as input in order to predict the sentiment labels thereof. Here, you will implement some data pre-processing and two major types of RNN “cell” (i.e. one time-step of computation). These are then used in other RNN modules that we provide to process entire sequences.

**Q1: Data processing** The reviews in the SST dataset come in various lengths. In the previous models we have looked at in the class, this has not been an issue because they rely either on a bag-of-words representation (Deep Averaging Network) or a fixed-sized window of previous tokens (Feed-Forward Language Model). RNNs, however, require the use of *padding*: given a batch of reviews of various lengths, we pad the shorter sequences with a special padding token so that all sequences are as long as the longest one. In `data.py`, please implement the `pad_batch` method. Please read the method signature and docstring carefully for details on the input and output. [5 pts]

**Q2: Vanilla RNN Cell** The “cell” of an RNN does one time-step of computation. For a Vanilla RNN, we saw that this was

$$h_t = \tanh(W_h h_{t-1} + b_h + W_x x_t + b_x)$$

where  $h_{t-1}$  is the previous hidden state,  $x_t$  is the current input, and the  $W$ s and  $b$ s are parameters for linear transformations.

In `model.py`, implement this computation in `VanillaRNNCell.forward`. The initializer defines the linear layers that you will need. [10 pts]

**Q3: LSTM Cell** An LSTM cell computes the next hidden state and *memory* based on the previous hidden state and memory together with the current input. Please consult these slides for the entire set of equations (and details about motivation).

In `model.py`, implement this computation in `LSTMCell.forward`. The initializer defines the linear layers that you will need. [15 pts]

## 3 Running the Classifier [15 pts]

`run.py` contains a basic training loop for SST classification, using the last hidden state of an RNN. It will record the training and dev loss at each epoch, and save the best model according to dev loss. At the end, it samples 10 random dev data points and prints the review, the gold label, and the model’s prediction.

**Q1: Four different runs** By default, a Vanilla RNN will be used. You can use an LSTM by specifying `--lstm` as a command-line argument. Following this paper, we have added dropout to the non-recurrent connections (i.e. from the inputs and to the output) of the model.

Please run each of the following variations. For each run, include in your `readme.pdf`: the best dev loss, the epoch at which the best dev loss was achieved, and the best model’s dev accuracy. [8 pts]

- Vanilla RNN, default parameters. (This is just `run.py` with no command-line arguments.)
- Vanilla RNN, with  $L_2$  regularization (via `--l2`) at  $1e-4$  and dropout (via `--dropout`) at 0.5.
- LSTM, default parameters.
- LSTM, with  $L_2$  regularization (via `--l2`) at  $1e-4$  and dropout (via `--dropout`) at 0.5.

**Q2: Inspecting outputs** For the fourth run above, please include in your `readme.pdf` the 10 random dev examples, with gold labels and model predictions here. In 2-3 sentences, describe what you see and observe any trends in what the model gets right and what (and/or how) it gets things wrong. [7 pts]

## 4 Testing your code

In the dropbox folder for this assignment, we will include a file `test_all.py` with a few very simple unit tests for the methods that you need to implement. You can verify that your code passes the tests by running `pytest` from your code's directory, with the course's conda environment activated.

## Submission Instructions

In your submission, include the following:

- `readme.(txt|pdf)` that includes your answers to §1 and §3.
- `hw6.tar.gz` containing:
  - `run_hw6.sh`. This should contain the code for activating the conda environment and your run commands for §3 above. You can use `run_hw2.sh` from the previous assignment as a template.
  - `data.py`
  - `model.py`