

Lecture 8: Recurrent Neural Networks

Madhavan Mukund and Pranabendu Misra

Advanced Machine Learning 2021


Dealing with sequences

- Conventional neural networks map single inputs to single outputs
 - Each input/output may be a vector of values
- These are *Feed Forward Networks*.





Dealing with sequences

- Conventional neural networks map single inputs to single outputs
 - Each input/output may be a vector of values
- These are *Feed Forward Networks*.
- Some classification tasks require mapping a sequence of inputs to an output
 - Identifying a music or video clip




Dealing with sequences

- Conventional neural networks map single inputs to single outputs
 - Each input/output may be a vector of values
- These are *Feed Forward Networks*.
- Some classification tasks require mapping a sequence of inputs to an output
 - Identifying a music of video clip
- Others require mapping a single input to a sequence of outputs
 - Generating a caption for an image





Dealing with sequences

- Mapping sequences to sequences
 - Language translation — read an entire input sentence, then generate output






Dealing with sequences

- Mapping sequences to sequences
 - Language translation — read an entire input sentence, then generate output
- Mapping sequences to sequences on the fly
 - Predict the next word in a sentence




Dealing with sequences

- Mapping sequences to sequences
 - Language translation — read an entire input sentence, then generate output
- Mapping sequences to sequences on the fly
 - Predict the next word in a sentence
- Context is important
 - The handwritten word is clearly **defence**
 - The **n** in isolation is illegible




Incorporating memory



- Input sequence $x^{(1)}, x^{(2)}, \dots, x^{(t)}$
- Output sequence $\hat{y}^{(1)}, \hat{y}^{(2)}, \dots, \hat{y}^{(t)}$
- Allow $\hat{y}^{(t)}$ to also depend on previous inputs $x^{(1)}, x^{(2)}, \dots, x^{(t-1)}$
- **Hidden state** : $h^{(t)}$
 - $h^{(t)}$ depends on current input and previous state
 - $h^{(t)} = f(W^{hx}x^{(t)} + W^{hh}h^{(t-1)} + b_h)$
- Output is a function of the current state
 - $y^{(t)} = g(W^{yh}h^{(t)} + b_y)$




Time unrolling





Time unrolling




Time unrolling



Time unrolling




Time Unrolling and Back Propagation Through Time




- Time Unrolling makes it a Feed-Forward Network
- So we can do back-propagation to update the weights

Vanishing and Exploding Gradients



- Unfortunately, we end-up with a very deep network
- Only the most recent parts of the input sequence are remembered; earlier parts are forgotten
- Back-Propagation suffers from vanishing or exploding gradients when unrolled over many time steps.

Vanishing and Exploding Gradients



- Truncated BPTT fixes some of these issues
- Unfortunately, long-term context is lost

Why this happens:

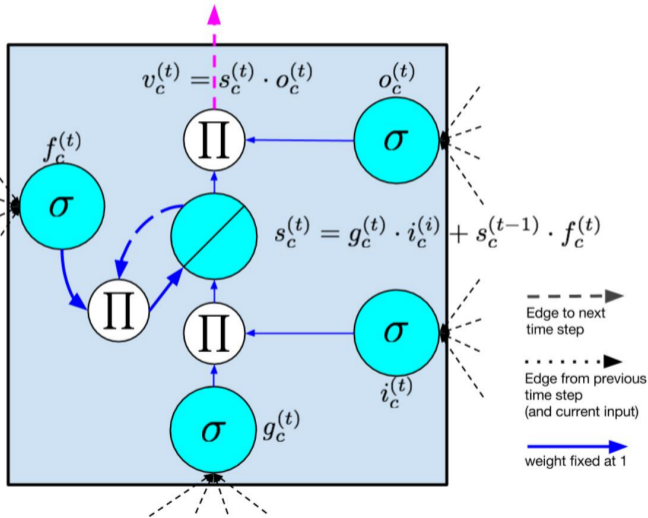
- Every piece x_i of the input sequence is treated in the same manner by the RNN.
- So important and non-important pieces both modify the hidden state, causing important pieces of the input sequences further in the past to be forgotten.

Why this happens:


- Every piece x_i of the input sequence is treated in the same manner by the RNN.
- So important and non-important pieces both modify the hidden state, causing important pieces of the input sequences further in the past to be forgotten.

A solution:


- One idea is to not update the hidden state all the time, but only when necessary.
Use *gates* to control information flow
- A gate is typically a 1-layer neural network with sigmoid activation, that takes the current input $x^{(t)}$ and the previous state $s^{(t)}$ and outputs a vector with values from $[0, 1]$.
0 means gate is closed; 1 means gate is open
- We thus arrive at *Gated RNNs*. Intuitively, gates learn to distinguish important and non-important information. They only let important information update the internal-state.



- Long Short Term Memory (LSTM) are a popular variant of gated RNNs.
- They use multiple gates to decide how the internal state $s_c^{(t)}$ is updated.
- Here, $x_c^{(t)}$ is the input at time t , which is supplied to all the gates along with the previous internal state $s_c^{(t-1)}$
- Π represents point-wise product of two vectors.



- $f_c^{(t)}$ is the *Forget Gate*. Decides what bits of $s_c^{(t-1)}$ is retained in $s_c^{(t)}$ and what is forgotten.
- $i_c^{(t)}$ is the *Input Gate*. Decides what bits of the input $x_c^{(t)}$ are added to $s_c^{(t)}$.
- $g_c^{(t)}$ is the *Input Node*, which is equivalent to an ordinary neuron. It's output is what is actually added to $s_c^{(t)}$ instead of $x_c^{(t)}$.



- We update the internal state as $s_c^{(t)} = \Pi(g_c^{(t)}, i_c^{(t)}) + \Pi(s_c^{(t-1)}, f_c^{(t)})$.
- The output gate $o_c^{(t)}$ decides what bits of the internal state $s_c^{(t)}$, is used for the output
- The output is $v_c^{(t)} = \Pi(s_c^{(t)}, o_c^{(t)})$

LSTM unrolled in time

