Deep Learning – NPFL114 – Exam, 17th January 2017

Write your answers to the following questions. When you are ready, notify me and we will go through the answers together, giving you a chance to fix eventual inaccuracies.

You can get a hint for any question, but you will get slightly less points for such questions.

1) Training Neural Network

Assume the artificial neural network on the right, with mean square error loss and gold output of 3. Compute the values of all weights w_i after performing a SGD update with learning rate 0.1.

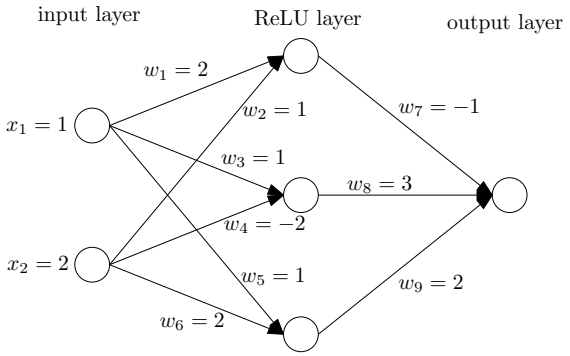

2) Training Neural Network

Assume the artificial neural network on the right, with negative log likelihood (cross-entropy) loss and gold distribution $(0, 0, 1)$, i.e., the last class is the gold one. Compute the values of all weights w_i after performing a SGD update with learning rate 0.1.

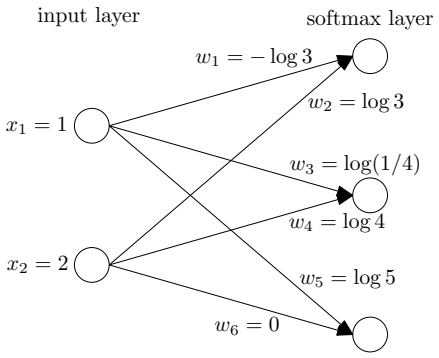

3) RMSProp

Write down the RMSProp algorithm.

4) Dropout

Describe the dropout method and explain how is it used during training and during inference.

5) Convolution

Specify how convolution of a given image is computed. Assume the input is an image I of size $H \times W$ with C channels, the kernel K has size $N \times M$, the stride is $T \times S$, the operation performed is in fact cross-correlation (as usual in convolutional neural networks) and that O output channels are computed. Spell out the computations for both *SAME* and *VALID* padding schemes.

6) GRU

Explain how the Gated Recurrent Unit cell operates.

7) Encoder-decoder with Attention

Draw/write how an encoder-decoder architecture is used for machine translation (both during training and during inference). Furthermore, elaborate on how attention mechanism works.

8) REINFORCE

Describe reinforcement learning abstraction, sketch policy gradient methods (notably, suggest an exemplary architecture for a policy network) and write down the REINFORCE algorithm.