

Back-to-School savings now on! Purchase before September 30th, and get 25% off Socrative Pro. 

Remind me later

Enter coupon code "BTS20" at checkout. [Learn More](#)

CONTI6128

Save and Exit

CAS Advanced - Day 7



Align Quiz to Standard

Enable Sharing
SOC-47868232

1. Why do people use encoder-decoder RNNs rather than plain sequence-to-sequence RNNs for automatic translation ?



- i* In general, if you translate a sentence one word at a time, the result will be terrible. It is much better to read the whole sentence first and then translate it. A plain sequence-to-sequence RNN would start translating a sentence immediately after reading the first word, while an encoder-decoder RNN will first read the whole sentence and then translate it.



2. How could you combine a convolutional NN with an RNN to classify videos ?



- i* To classify videos based on the visual content, one possible architecture could be to take one frame per second, then run each frame through a convolutional NN, feed the output of the CNN to a sequence-to-vector RNN, and finally run its output through a softmax layer, giving you all the class probabilities. For training, you would just use cross entropy as the cost function. If you wanted to use the audio for classification as well, you could



convert every second of audio to a spectrograph, feed this spectrograph to a CNN, and feed the output of this CNN to the RNN.

3. How can you deal with variable-length input sequences ?

- i The simplest option is to set the `sequence_length` parameter function when calling `static_rnn()` or `dynamic_rnn()` in Tensorflow. Another option is to pad the smaller inputs (with zeros) to make them the same size as the largest inputs.



4. How can you deal with variable-length output sequences ?

- i If you know in advance the length of each output sequence, you can use the `sequence_length` parameter. If you don't know in advance the length, you can use the padding trick : always output the same size sequence, but ignore any outputs that come after the end-of-sequence token.



Add a Question

Multiple Choice

True / False

Short Answer