

Data Analysis and Machine Learning: Recurrent neural networks

Morten Hjorth-Jensen^{1,2}

¹Department of Physics, University of Oslo

²Department of Physics and Astronomy and National Superconducting Cyclotron Laboratory, Michigan State University

Mar 11, 2019

Recurrent neural networks: Overarching view

Till now our focus has been, including convolutional neural networks as well, on feedforward neural networks. The output or the activations flow only in one direction, from the input layer to the output layer.

A recurrent neural network (RNN) looks very much like a feedforward neural network, except that it also has connections pointing backward.

RNNs are used to analyze time series data such as stock prices, and tell you when to buy or sell. In autonomous driving systems, they can anticipate car trajectories and help avoid accidents. More generally, they can work on sequences of arbitrary lengths, rather than on fixed-sized inputs like all the nets we have discussed so far. For example, they can take sentences, documents, or audio samples as input, making them extremely useful for natural language processing systems such as automatic translation and speech-to-text.

Set up of an RNN

The figure here displays a simple example of an RNN, with inputs x_t at a given time t and outputs y_t . Introducing time as a variable offers an intuitive way of understanding these networks. In addition to the inputs x_t , the layer at a time t receives also as input the output from the previous layer $t - 1$, that is y_{t-1} .

This means also that we need to have weights that link both the inputs x_t to the outputs y_t as well as weights that link the output from the previous time y_{t-1} and y_t . The figure here shows an example of a simple RNN.