Linear Autoencoder Pretraining for Orthogonal RNN with Linear Memory Networks

Antonio Carta

July 30, 2019

Abstract

Orthogonal recurrent neural networks solve the vanishing gradient problem by parameterizing the recurrent connections using an orthogonal matrix. We propose a novel architecture, called Linear Memory Network, made of a linear recurrence and a separate nonlinear component. The network is equivalent to an Elman RNN but due to the linear recurrence, it is less sensitive to vanishing gradient problems. We devise a pretraining schema that initializes the network exploiting a linear recurrent autoencoder. The optimal solution of a linear authoencoder is an orthogonal RNN which is able to memorize the entire input sequence into its hidden state. We argue that our approach is superior to a random orthogonal initialization due to the more efficient encoding. Experimental results show that our approach outperforms orthogonal models with random orthogonal initialization and LSTM on sequential MNIST and permuted MNIST.