

May 06, 2016

Alexander Koller Sebastian Stober

*with figures from deeplearningbook.org







[Y. Bengio and Y. Lecun, 1995]

Convolution & Recurrence



Recurrence: Motivation

- process sequential data
- capture history of inputs/states
- share parameters through a very deep computational graph
 - output is a function of the previous output
 - produced using the same update rule applied to the previous outputs.
- different from convolution across time steps

The general idea



Convolution & Recurrence

Universitä,

ro rsdam Deep Learning

for NLP



Cyclic Connections

- computational graph includes cycles (recursion)
- represent influence of the present value of a variable on its own value at future time steps

- unfolding to yield a graph that does not involve recurrence
 - => gets very deep very quickly



- regardless of the sequence length, the learned model always has the same input dimensionality
- 2. can use the same transition function f with the same parameters at every time step



- gradient computation for unfolded loss function w.r.t parameters very expensive
- O(T) where T is history length
- no parallelization (sequential dependence)



Hidden State

- h(t) as a kind of "lossy summary" of the taskrelevant aspects of the history up to t
 - lossy compression necessary
 - selectivity based on training criterion (cost)

 most demanding situation: rich enough representation h(t) to allow approximate recovery of input sequences (autoencoder)

Output for entire sequence



Universitäx

ro rsdam Deep Learning

for NLP

Output at every step



Such a recurrent network of with finite size can compute any function computable by a Turing machine.

Convolution & Recurrence

Universitä

• Porsdam

Deep Learning

for NLP

Bi-Directional RNN



Universita

Deep Learning

for NLP

- h(t) relevant summary of past (forward)
- g(t) relevant summary of future (backward)
- extendable to 2D inputs
- restriction: samelength sequences

Image labeling Deep Learning



Convolution & Recurrence

Universitäx

for NLP

2016-04-29 12



Encoder-Decoder Sequence-to-Sequence

- encoder (reader)
 - read input sequence
 - generate hidden state



- generate output sequence from hidden state
- variable length



Recursive NNs



Convolution & Recurrence

Universitä,

• Porsidam

Deep Learning

for NLP



RNNs: Challenges

- Training can be slower than for feedforward or convolutional networks.
- Vanishing/exploding gradients.
- Difficulties in learning to use hidden state to remember information about the distant past.
 - Extremely important in NLP!



Gradient Clipping

Without clipping

With clipping







Long Short-Term Memory; Hochreiter & Schmidhuber 97

Convolution & Recurrence





Figure 1: Illustration of (a) LSTM and (b) gated recurrent units. (a) i, f and o are the input, forget and output gates, respectively. c and \tilde{c} denote the memory cell and the new memory cell content. (b) r and z are the reset and update gates, and h and \tilde{h} are the activation and the candidate activation.

Gated Recurrent Units; Cho et al. 2014

2016-04-29 18

Convolution & Recurrence





- Recurrent neural networks extremely useful for modeling natural language.
- Special challenges that need to be overcome.
- Many possible solutions available; try them!

Convolution & Recurrence