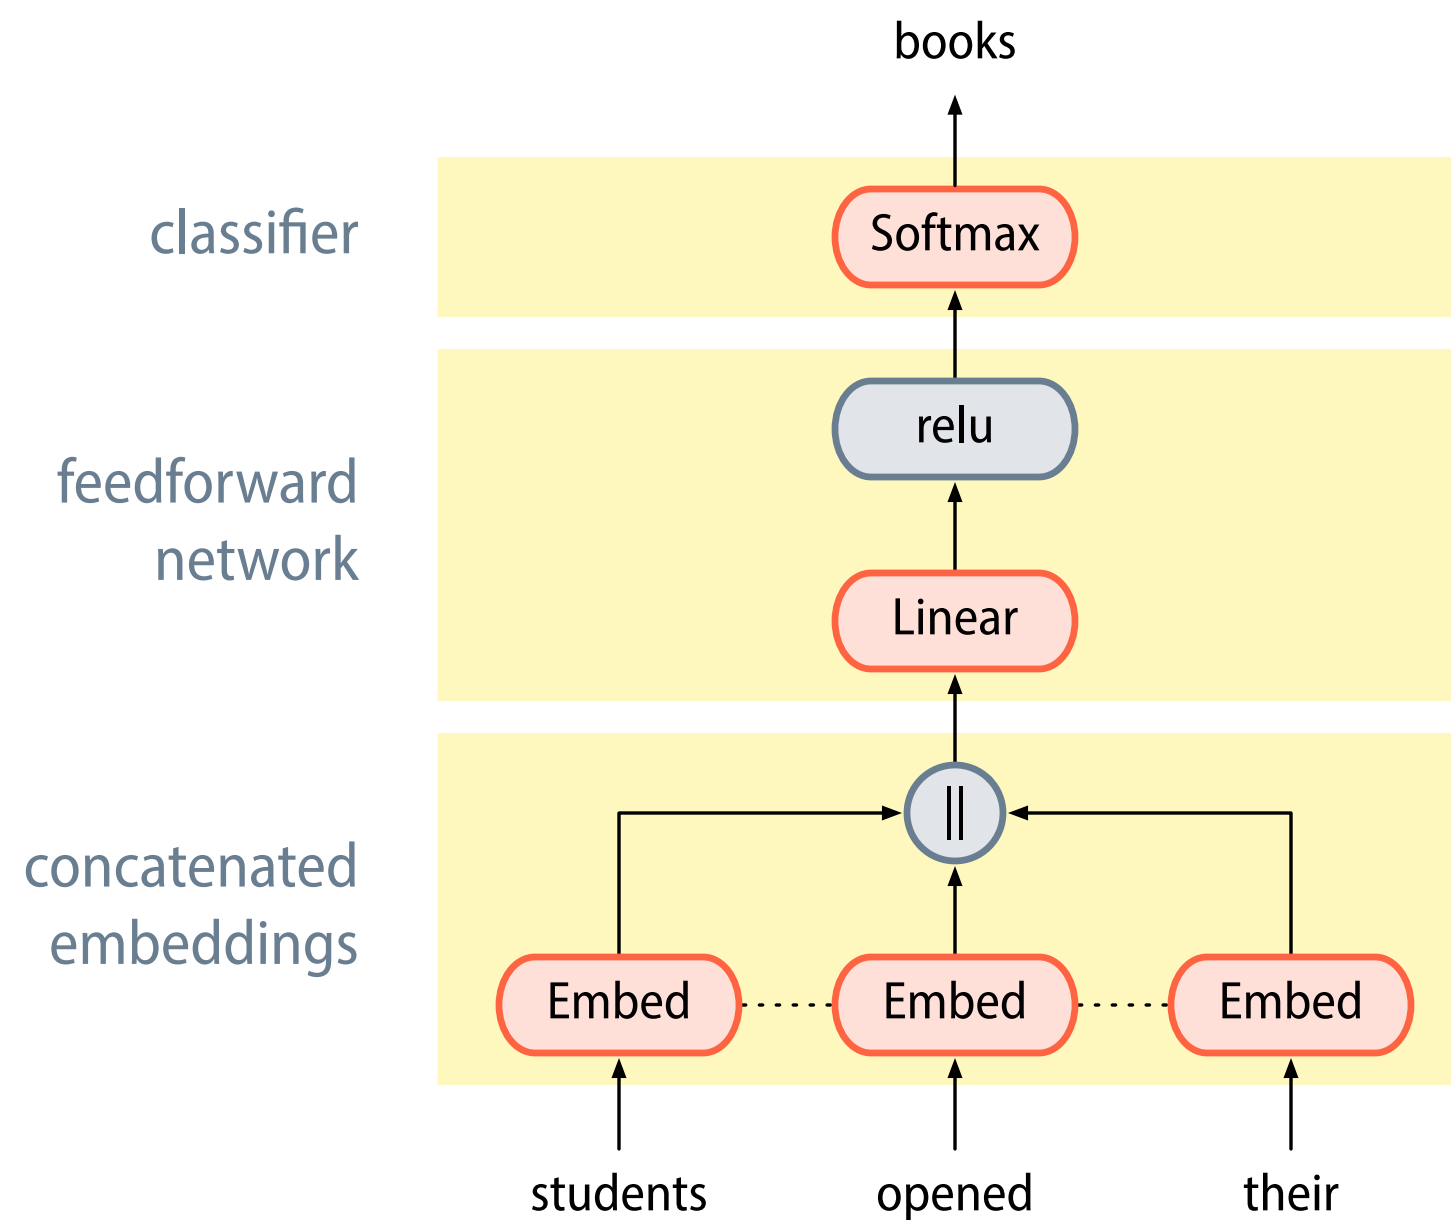


# Recurrent neural network language models

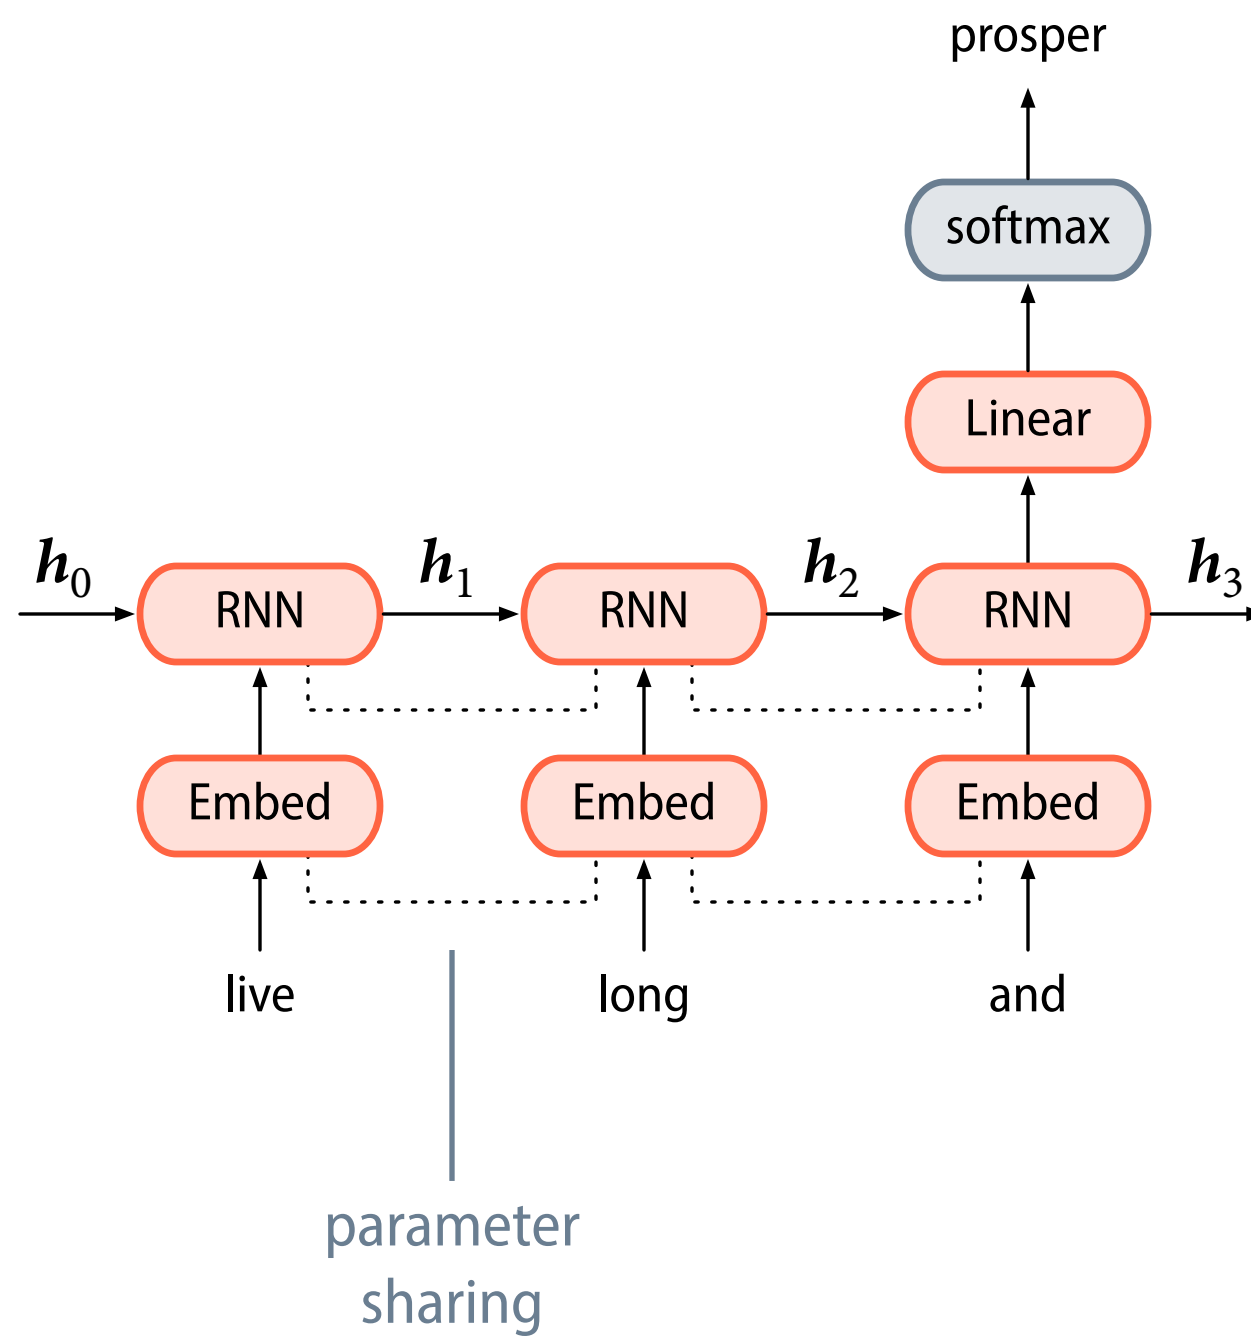
Marco Kuhlmann

Department of Computer and Information Science

# Reminder: Fixed-window neural language model



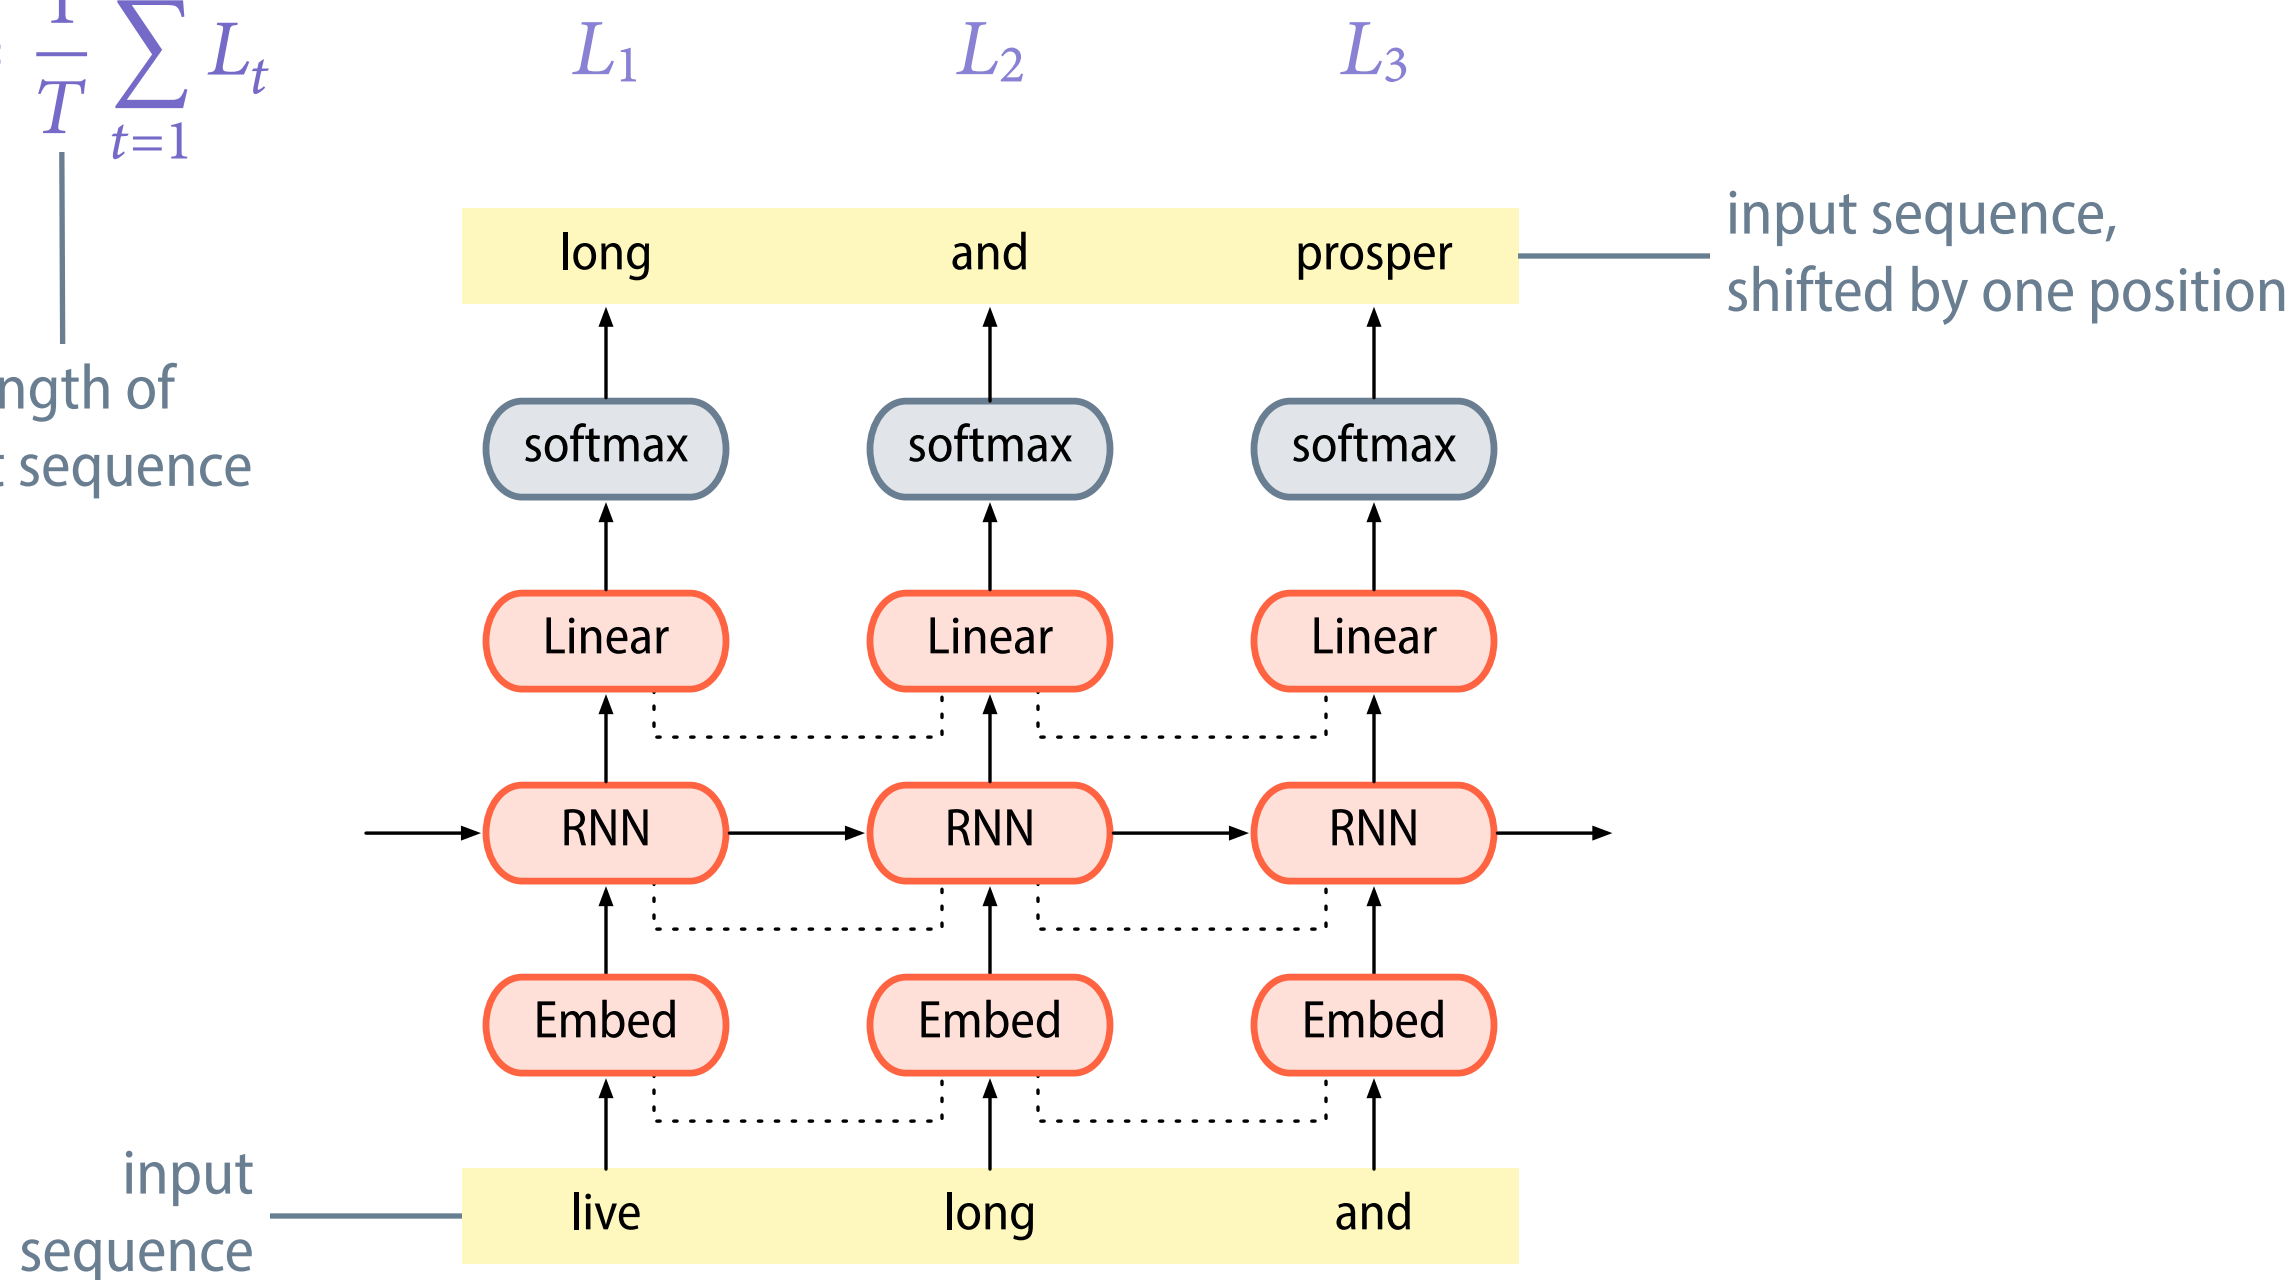
# RNN language model



# Training RNN language models

$$L = \frac{1}{T} \sum_{t=1}^T L_t$$

length of input sequence




# Practical issues when training RNNs

- In practice, backpropagation through time is often truncated after some fixed number of steps.

corresponds to truncating the input sequence

- When creating batches of sequences with unequal lengths, shorter sequences need to be padded.
- To minimize padding, we can sort sequences by length, and try to create minibatches in which sequences have the same length.

# Recurrent neural networks for sequence labelling



'I only want to live in peace, plant potatoes, and dream!' – Moomin