

Introduction to Deep Learning

CNNs Part II

Andreas Krug, M.Sc.

ankrug@uni-potsdam.de

5 -

05. November 2018



Group exercise

Receptive field



Receptive field

1D-conv, kernel=3

1D-conv, kernel=3



1 pixel sees 3 their receptive fields overlap by 2 receptive field = 3*3 - (3-1)*2 = 5

1 pixel sees 3 receptive field = 3







[Y. Bengio and Y. Lecun, 1995]

CNN from Y. LeCun and Y. Bengio: **Convolutional Networks for Images, Speech, and Time-Series**, in Arbib, M. A. (Eds), *The Handbook of Brain Theory and Neural Networks, MIT Press, 1995*.





What is the:



- number of parameters in each layer
- total number of parameters in the CNN
- number of parameters in a single fully-connected layer

?

• size of the receptive field in each layer













Layer 2: 2x2 pooling => 6x6-tiles





Layer 2: 2x2 pooling => 6x6-tiles Layer 2: 2x subsampling => 12x12 6x6-tiles at stride 2





Layer 2: 2x2 pooling => 6x6-tiles Layer 2: 2x subsampling => 12x12 6x6-tiles at stride 2

Layer 3: 5x5 filters => 5x5 6x6-tiles at stride 2





Layer 2: 2x2 pooling => 6x6-tiles Layer 2: 2x subsampling => 12x12 6x6-tiles at stride 2

Layer 3: 5x5 filters => 5x5 6x6-tiles at stride 2 => 8x8 14x14-tiles at stride 2





Layer 2: 2x2 pooling => 6x6-tiles Layer 2: 2x subsampling => 12x12 6x6-tiles at stride 2

Layer 3: 5x5 filters => 5x5 6x6-tiles at stride 2 => 8x8 14x14-tiles at stride 2

Layer 4: 2x2 pooling => 16x16-tiles





Layer 2: 2x2 pooling => 6x6-tiles Layer 2: 2x subsampling => 12x12 6x6-tiles at stride 2

Layer 3: 5x5 filters => 5x5 6x6-tiles at stride 2 = 8x8 14x14-tiles at stride 2

Layer 4: 2x2 pooling => 16x16-tiles Layer 4: 2x subsampling => 4x4 16x16-tiles at stride 4





Layer 2: 2x2 pooling => 6x6-tiles Layer 2: 2x subsampling => 12x12 6x6-tiles at stride 2

Layer 3: 5x5 filters => 5x5 6x6-tiles at stride 2 => 8x8 14x14-tiles at stride 2

Layer 4: 2x2 pooling => 16x16-tiles Layer 4: 2x subsampling => 4x4 16x16-tiles at stride 4

Layer 5: 4x4 filters => 1x1 28x28-tile



Assignments until next week

- Recap: Bjarne & Jonathan
- Reading: CNN papers in detail Doodle for picking one out of six
- Programming exercise: High-level Tensorflow with tf.Estimator

Slides & assignments on: <u>https://mlcogup.github.io/idl_ws18/</u>