Convolutional Neural Networks, Part I.

RECAP: 29/10

Rebekah Cramerus & Vageesh Saxena

MLPs

- Every element of a previous layer, is connected to every elements of the next layer.
- Feed-forward Artificial neural networks.
- 3 layers – I/P, O/P, and the hidden layers.
- Except for the I/P nodes, each node uses a nonlinear activation function.
- Uses Back-propagation for training.
- It can distinguish data that is not linearly separable

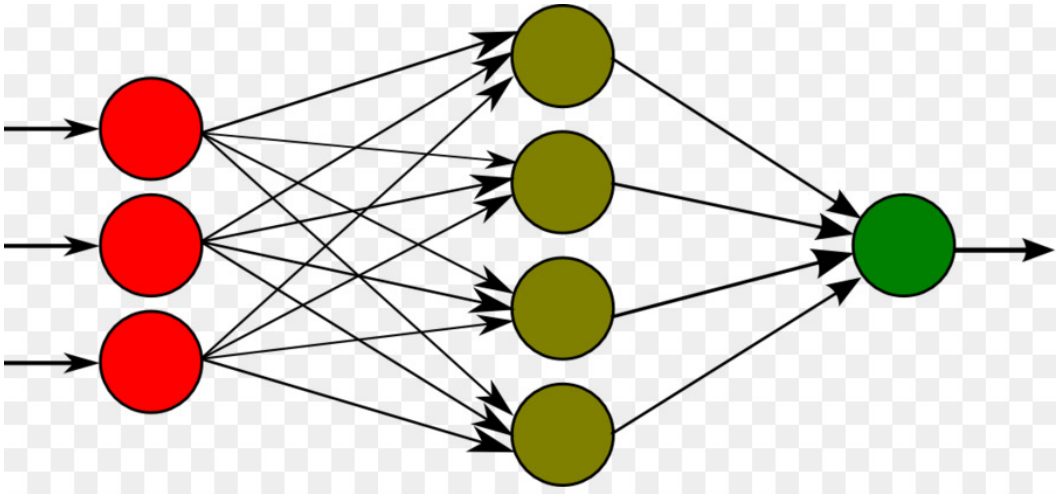

CNNs

- Deep, feed-forward network that uses mathematical operations called **convolutions** in place of general matrix multiplication.
- Useful for time series data and image data (with which it has had great success).

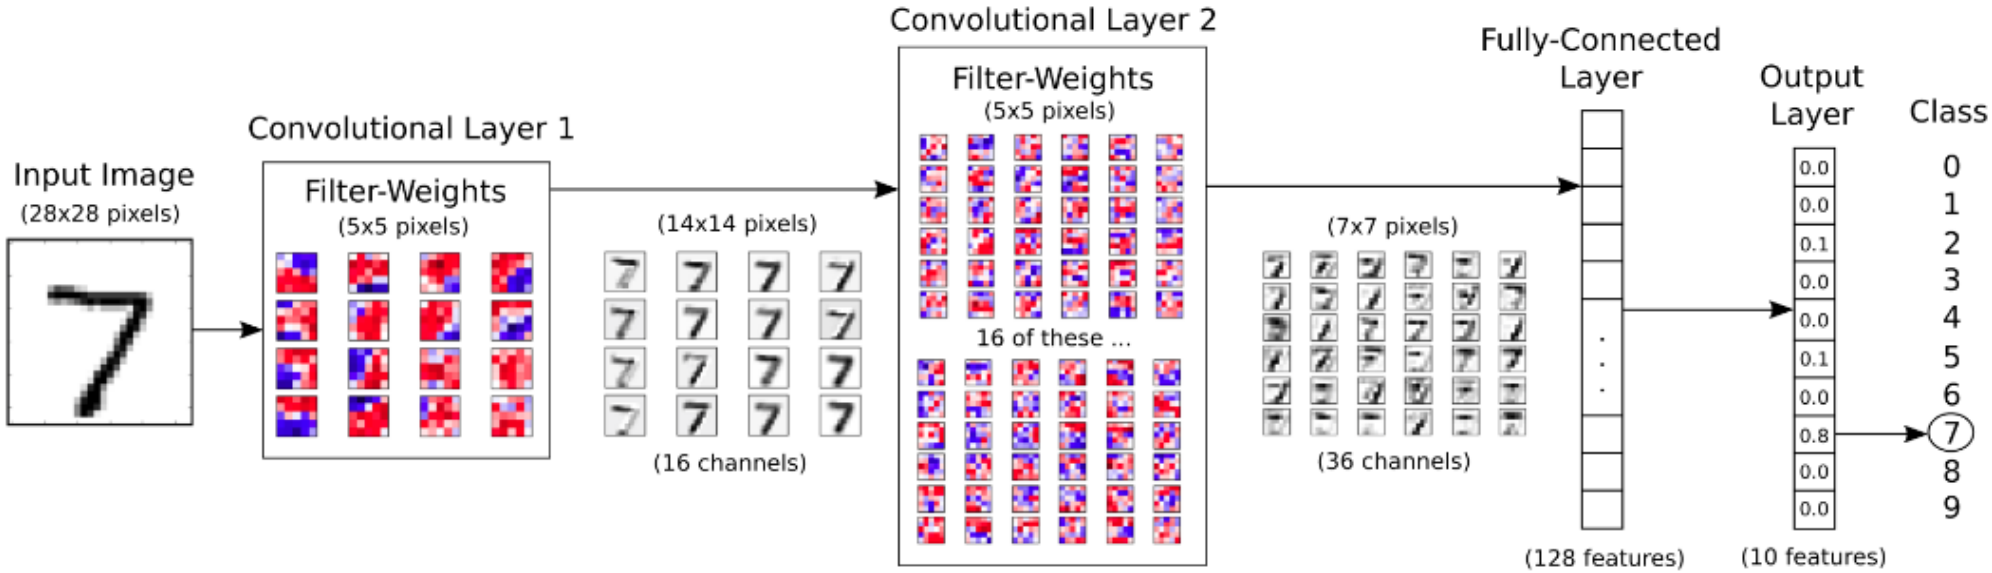
Input Image
(28x28 pixels)
Convolutional Layer 1
Filter-Weights
(5x5 pixels)
(14x14 pixels)
(16 channels)
Convolutional Layer 2
Filter-Weights
(5x5 pixels)
16 of these ...
(7x7 pixels)
(36 channels)
Fully-Connected Layer
(128 features)
Output Layer
0.0
0.0
0.1
0.0
0.0
0.1
0.0
0.8
0.0
0.0
(10 features)
Class
0
1
2
3
4
5
6
7
8
9

Convolution

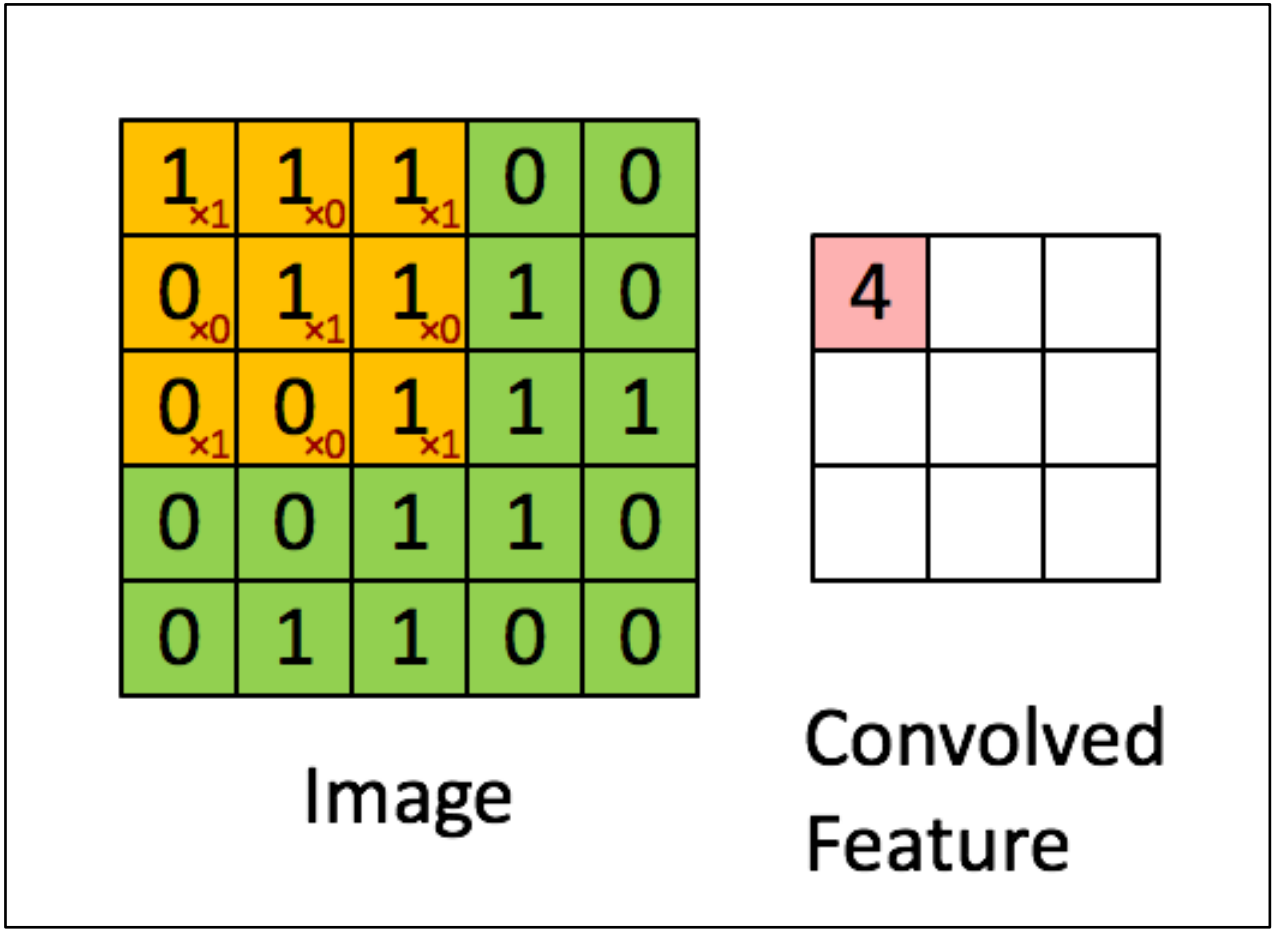

Parameter sharing

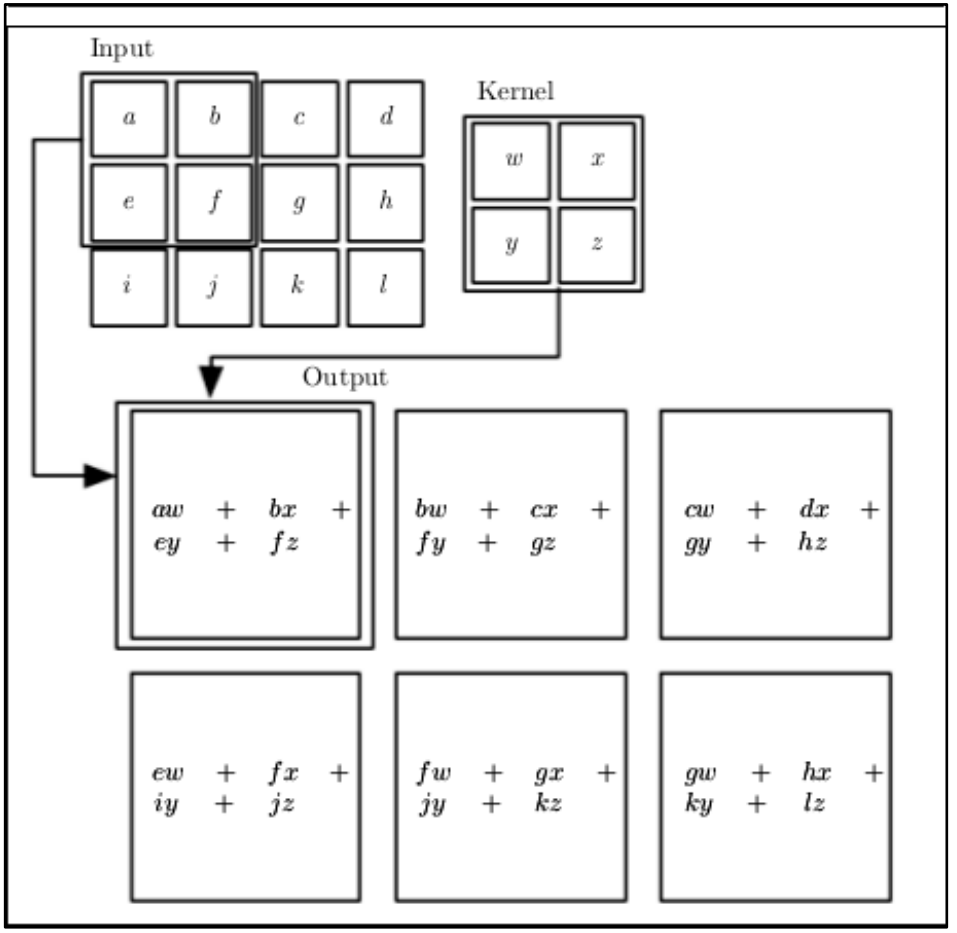

Sparse Connectivity

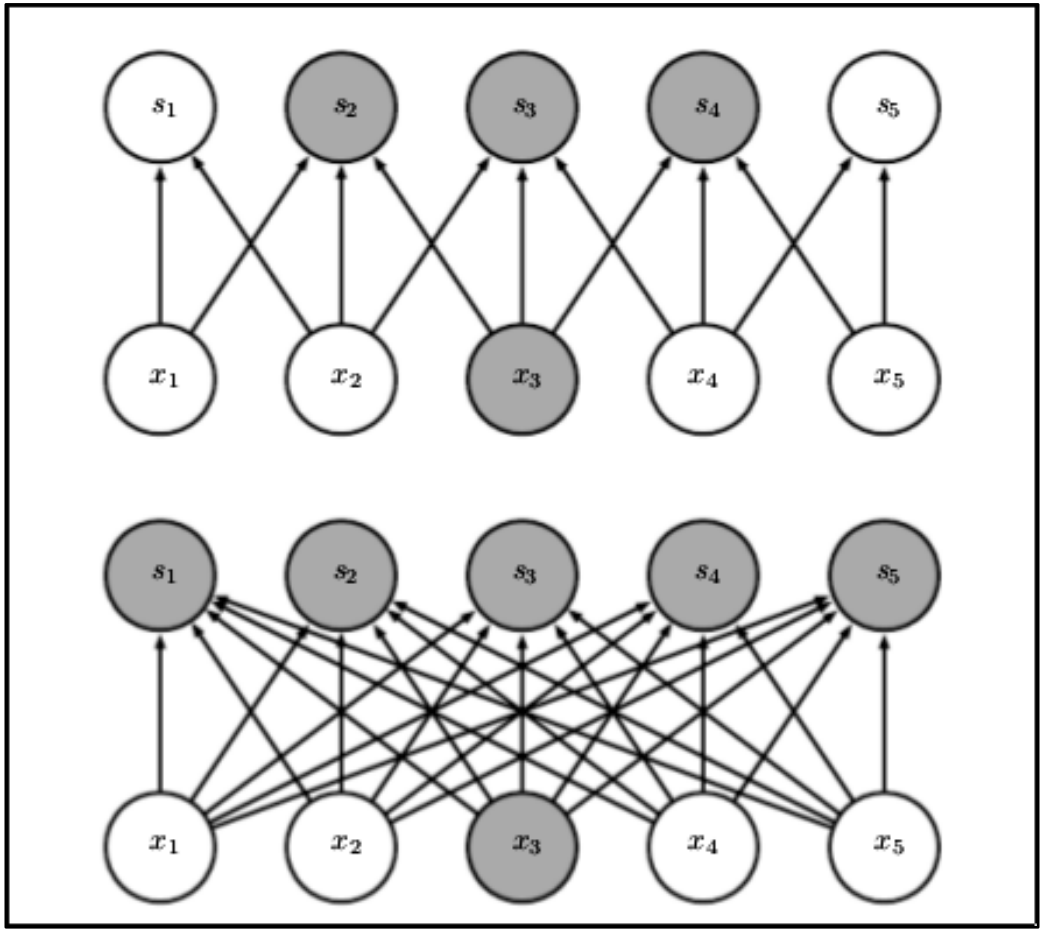

Pooling

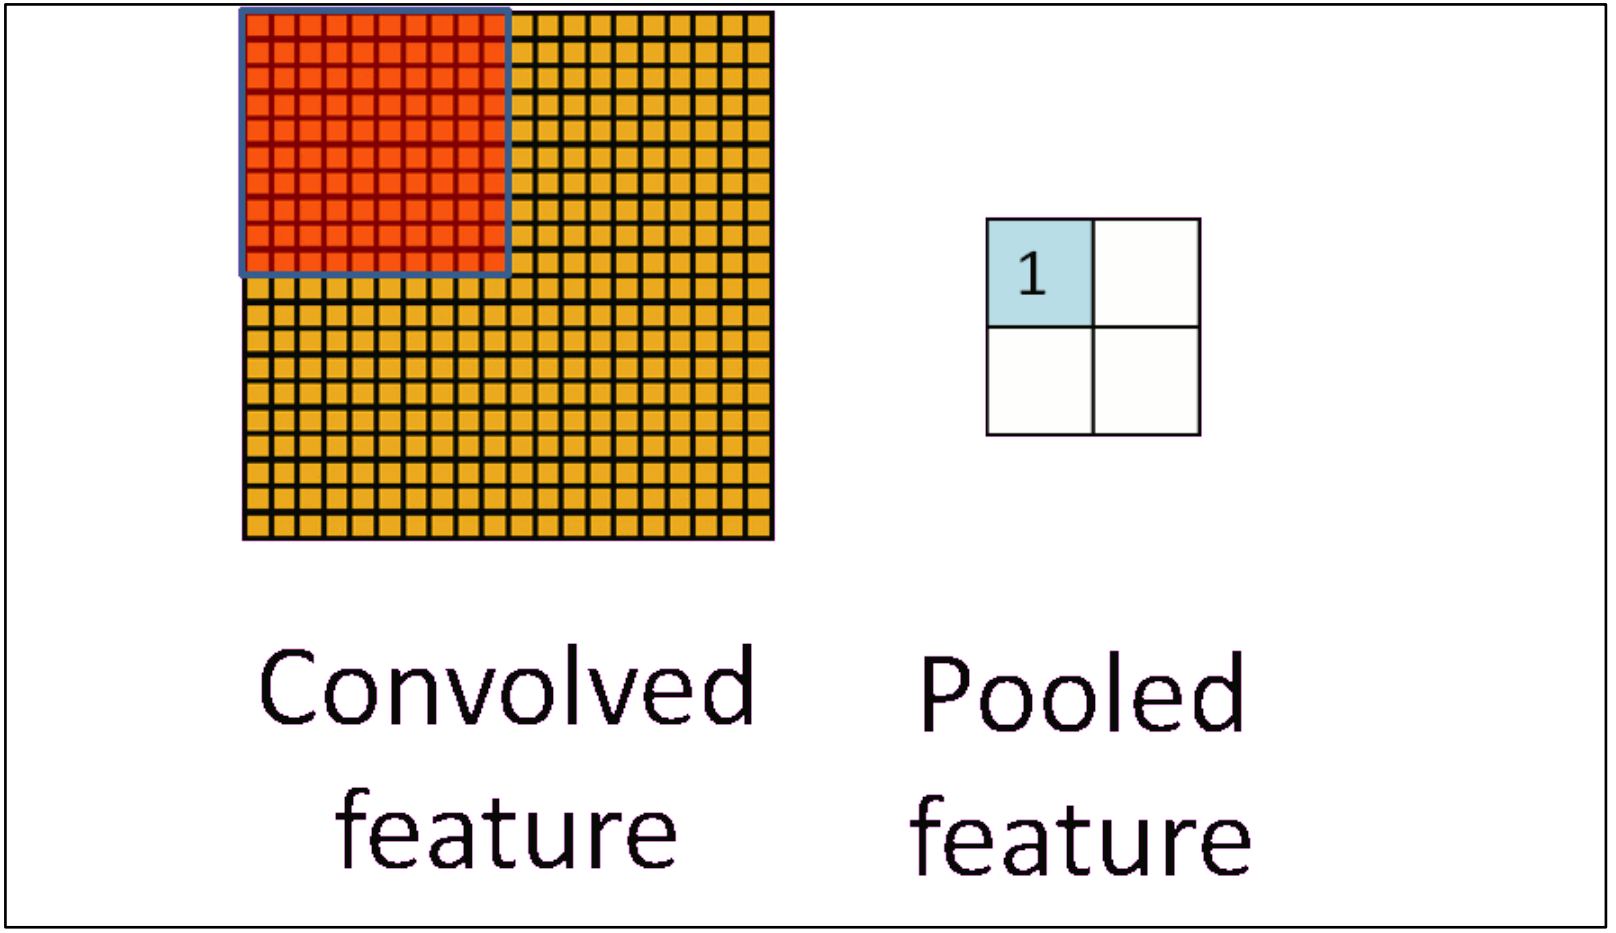

Stride

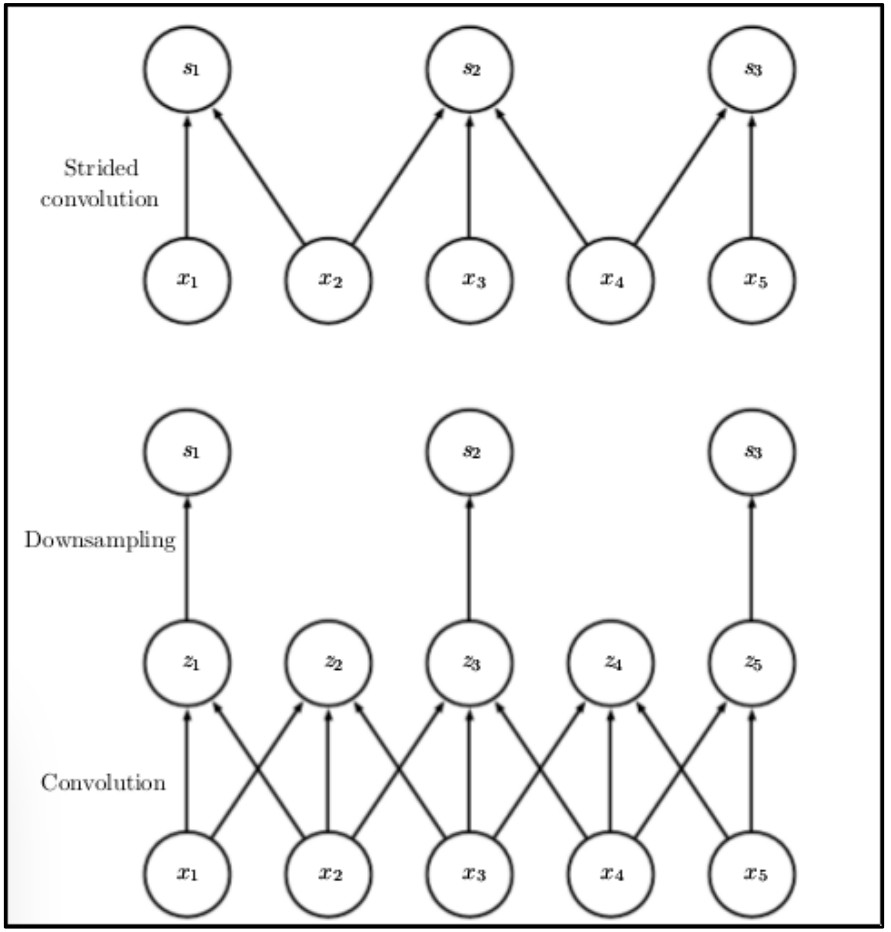
s_1
s_2
s_3
Strided convolution
x_1
x_2
x_3
x_4
x_5
s_1
s_2
s_3
Downsampling
z_1
z_2
z_3
z_4
z_5
Convolution
x_1
x_2
x_3
x_4
x_5

Padding: Valid vs. Same