

# Deep Learning

Using a Convolutional Neural Network

**Dr. – Ing. Morris Riedel**

Adjunct Associated Professor

School of Engineering and Natural Sciences, University of Iceland

Research Group Leader, Juelich Supercomputing Centre, Germany

LECTURE 2

## Convolutional Neural Networks & Tools

November 30<sup>th</sup>, 2017

Ghent, Belgium




UNIVERSITY OF ICELAND  
SCHOOL OF ENGINEERING AND NATURAL SCIENCES

FACULTY OF INDUSTRIAL ENGINEERING,  
MECHANICAL ENGINEERING AND COMPUTER SCIENCE



# Outline



# Outline of the Course

1. Deep Learning Fundamentals & GPGPUs
2. Convolutional Neural Networks & Tools
3. Convolutional Neural Network Applications
4. Convolutional Neural Network Challenges
5. Transfer Learning Technique
6. Other Deep Learning Models & Summary




# Outline

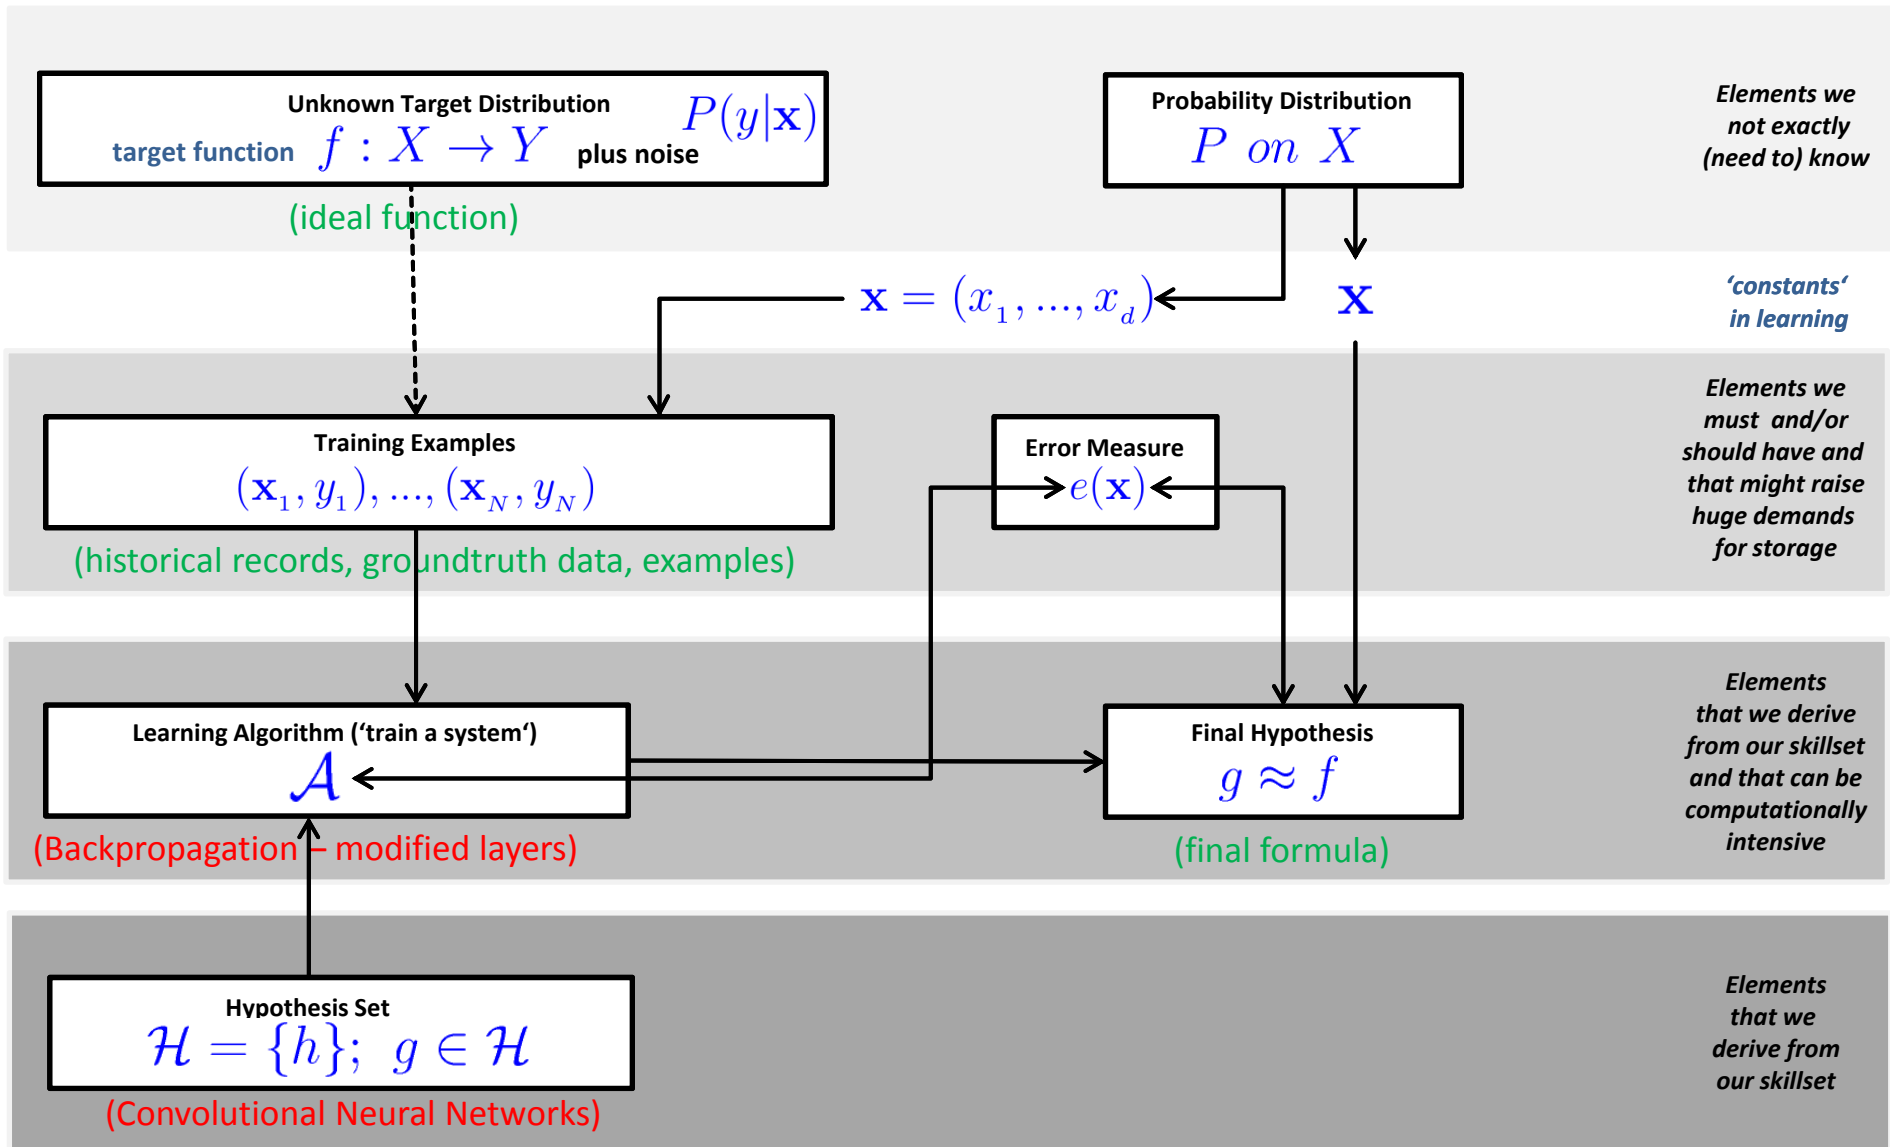
- Convolutional Neural Networks (CNNs)
  - Basic Principles & MNIST Application Example
  - Local Receptive Fields & Sliding
  - Revisit ANN Overfitting & Weight Problem
  - Shared Weights & Feature Maps
  - Advanced Application Examples
- Deep Learning Toolset
  - Timeline of Selected Relevant Tools
  - Low-level Deep Learning Libraries
  - Tensorflow, Caffe and Theano
  - Tensorflow Computational Graph
  - What is a Tensor & Chain Rule



# Convolutional Neural Networks (CNNs)



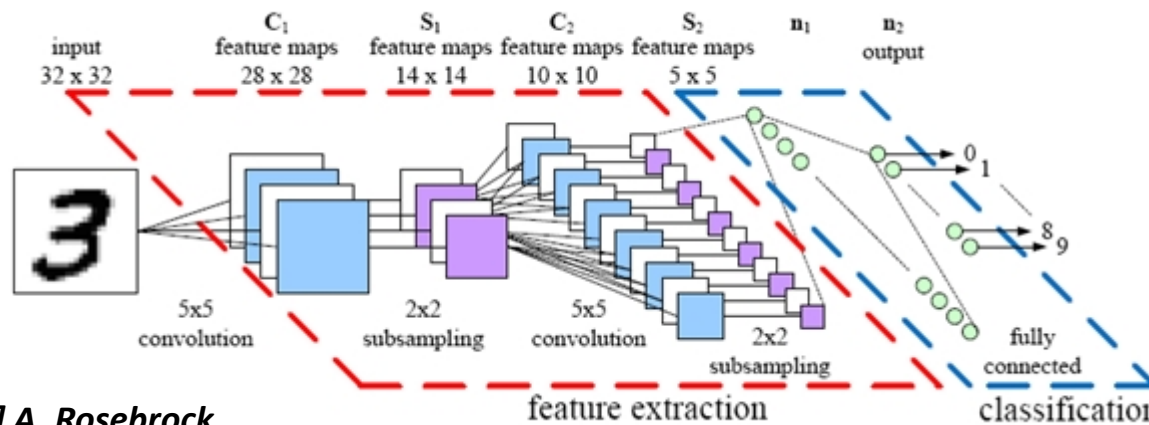

# Solution Tools: Convolutional Networks Learning Model



# CNNs – Basic Principles

- Convolutional Neural Networks (CNNs/ConvNets) implement a connectivity pattern between neurons inspired by the animal visual cortex and use several types of layers (convolution, pooling)
- CNN key principles are local receptive fields, shared weights, and pooling (or down/sub-sampling)
- CNNs are optimized to take advantage of the spatial structure of the data

- Simple application example
  - MNIST database written characters
  - Use CNN architecture with different layers
  - Goal: automatic classification of characters




[3] A. Rosebrock

[1] M. Nielsen

# CNNs – Principle Local Receptive Fields

- MNIST dataset example
  - 28 \* 28 pixels modeled as square of neurons in a convolutional net
  - Values correspond to the 28 \* 28 pixel intensities as inputs


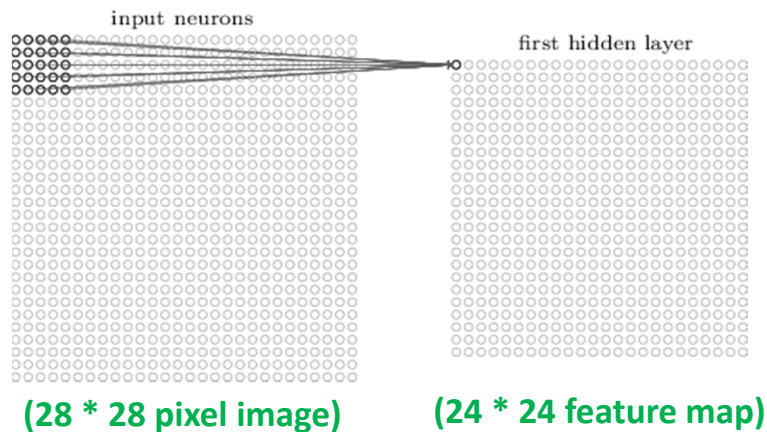


[1] M. Nielsen



# CNNs – Principle Local Receptive Fields & Sliding

- MNIST database example
  - Apply stride length = 1
  - Different configurations possible and depends on application goals
  - Creates 'feature map' of 24 \* 24 neurons (hidden layer)




[1] M. Nielsen

# CNNs –Example with an ANN with risk of Overfitting

- MNIST database example

- CNN: e.g. 20 feature maps with  $5 * 5$  (+bias) = **520 weights to learn**
- Apply ANN that is fully connected between neurons
- ANN: fully connected first layer with  $28 * 28 = 784$  input neurons
- ANN: e.g. 15 hidden neurons with  $784 * 15 =$  **11760 weights to learn**



(eventually lead to overfitting and much computing time)

[1] M. Nielsen

# Exercises – ANN Example Revisited – Count Parameters




# CNNs – Principle Shared Weights & Feature Maps

- Approach
  - CNNs use same shared weights for each of the  $24 \times 24$  hidden neurons
  - Goals: significant reduction of number of parameters (prevent overfitting)
  - Example:  $5 \times 5$  receptive field  $\rightarrow$  25 shared weights + shared bias

- Feature Map

- Detects one local feature
- E.g. 3: each feature map is defined by a set of  $5 \times 5$  shared weights and a single shared bias leading to  $24 \times 24$
- Goal: The network can now detect 3 different kind of features (many more in practice)
- Benefit: learned feature being detectable across the entire image





(shared weights are also known to define a kernel or filter)

[1] M. Nielsen

# CNNs – Principle of Pooling

- ‘Downsampling’ Approach
  - Usually applied directly after convolutional layers
  - Idea is to simplify the information in the output from the convolution
  - Take each feature map output from the convolutional layer and **generate a condensed feature map**
  - E.g. Pooling with  $2 \times 2$  neurons using ‘max-pooling’
  - Max-Pooling outputs the maximum activation in the  $2 \times 2$  region

hidden neurons (output from feature map)

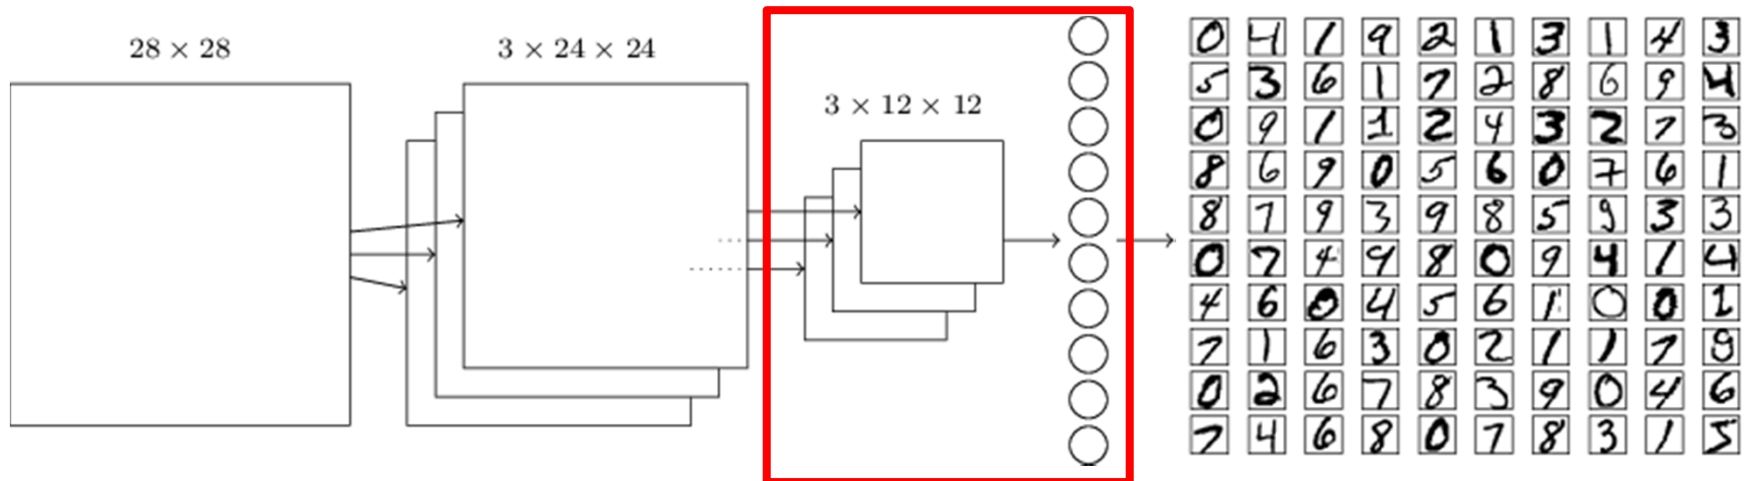


[1] M. Nielsen


# CNN – Application Example MNIST

- MNIST database example

- Full CNN with the addition of output neurons per class of digits
  - Apply ‘fully connected layer’: layer connects every neuron from the max-pooling outcome layer to every neuron of the 10 out neurons
  - Train with **backpropagation algorithm (gradient descent)**, only small modifications for new layers



- Approach works, except for **some bad training and test examples**



(another indicator that even with cutting edge technology machine learning never achieves 100% performance)

[1] M. Nielsen

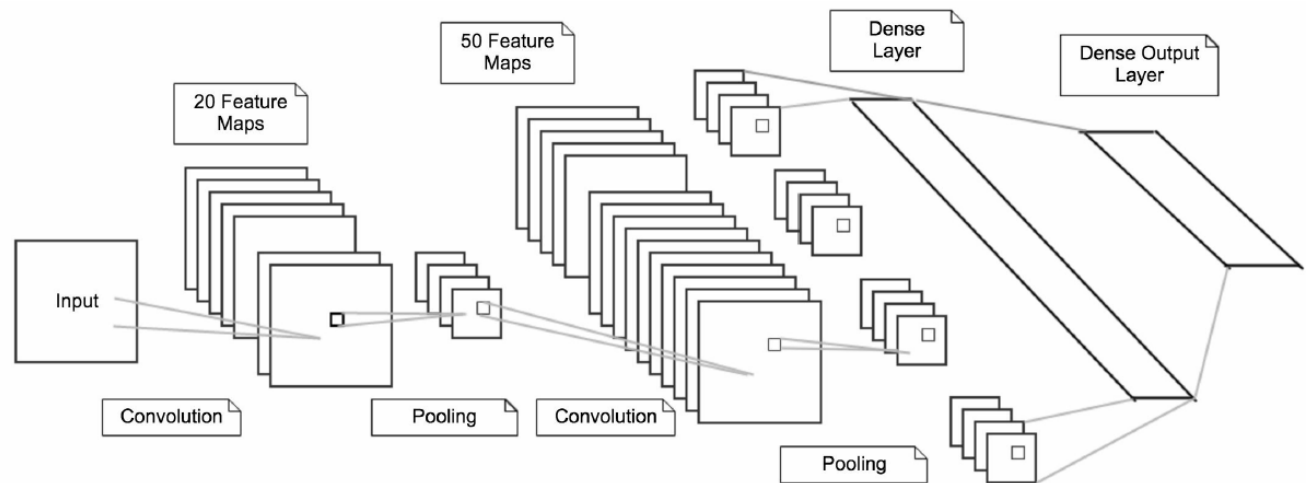
# Exercises – MNIST Dataset – CNN Model Example



# MNIST Dataset – CNN Model

```
from keras.datasets import mnist
from keras.models import Sequential
from keras.layers.core import Dense, Activation, Flatten
from keras.utils import np_utils
from keras import backend as K
from keras.layers.convolutional import Convolution2D, MaxPooling2D
from keras.optimizers import SGD, RMSprop, Adam

# model
class CNN:
    @staticmethod
    def build(input_shape, classes):
        model = Sequential()
        model.add(Convolution2D(20, kernel_size=5, padding="same", input_shape=input_shape))
        model.add(Activation("relu"))
        model.add(MaxPooling2D(pool_size=(2,2), strides=(2,2)))
        model.add(Convolution2D(50, kernel_size=5, border_mode="same"))
        model.add(Activation("relu"))
        model.add(MaxPooling2D(pool_size=(2,2), strides=(2,2)))
        model.add(Flatten())
        model.add(Dense(500))
        model.add(Activation("relu"))
        model.add(Dense(classes))
        model.add(Activation("softmax"))
        return model
```



[9] A. Gulli et al.



# MNIST Dataset – CNN Python Script

```
# parameters
NB_CLASSES = 10
NB_EPOCH = 20
BATCH_SIZE = 128
VERBOSE = 1
OPTIMIZER = 'Adam'
VALIDATION_SPLIT = 0.2
IMG_ROWS, IMG_COLS = 28, 28
INPUT_SHAPE = (1, IMG_ROWS, IMG_COLS)
```

```
# dataset 28 x 28 pixels
(X_train, y_train), (X_test, y_test) = mnist.load_data()
K.set_image_dim_ordering("th")
X_train = X_train.astype('float32')
X_test = X_test.astype('float32')
```

```
# normalization
X_train /= 255
X_test /= 255
```

```
# input convnet
X_train = X_train[:, np.newaxis, :, :]
X_test = X_test[:, np.newaxis, :, :]
```

```
# data output
print(X_train.shape[0], 'train samples')
print(X_test.shape[0], 'test samples')
```

```
# convert vectors to binary matrices of classes
Y_train = np_utils.to_categorical(y_train, NB_CLASSES)
Y_test = np_utils.to_categorical(y_test, NB_CLASSES)
```

```
# Simple CNN model
model = CNN.build(input_shape=INPUT_SHAPE, classes=NB_CLASSES)
```

```
# Compilation
model.compile(loss='categorical_crossentropy', optimizer=OPTIMIZER, metrics=['accuracy'])
```

```
# Fit the model
history = model.fit(X_train, Y_train, batch_size=BATCH_SIZE, epochs=NB_EPOCH, verbose=VERBOSE, validation_split=VALIDATION_SPLIT)
```

```
# evaluation
score = model.evaluate(X_test, Y_test, verbose=VERBOSE)
print('Test score:', score[0])
print('Test accuracy:', score[1])
```

- **OPTIMIZER: Adam** - advanced optimization technique that includes the concept of a momentum (a certain velocity component) in addition to the acceleration component of Stochastic Gradient Descent (SGD)
- Adam computes individual adaptive learning rates for different parameters from estimates of first and second moments of the gradients
- Adam enables faster convergence at the cost of more computation and is currently recommended as the default algorithm to use (or SGD + Nesterov Momentum)

*[12] D. Kingma et al.,  
'Adam: A Method for  
Stochastic Optimization'*

# MNIST Dataset – CNN Model – Output

```
[vsc42544@gligar01 deeplearning]$ head KERAS_MNIST_CNN.o1179880
60000 train samples
10000 test samples
Train on 48000 samples, validate on 12000 samples
Epoch 1/20
```

```
128/48000 [.....] - ETA: 10:06 - loss: 2.2997 - acc: 0.1250
256/48000 [.....] - ETA: 7:46 - loss: 2.2578 - acc: 0.1992
384/48000 [.....] - ETA: 6:58 - loss: 2.2127 - acc: 0.2083
512/48000 [.....] - ETA: 6:35 - loss: 2.1632 - acc: 0.2598
640/48000 [.....] - ETA: 6:20 - loss: 2.0934 - acc: 0.3234
```

```
[vsc42544@gligar01 deeplearning]$ tail KERAS_MNIST_CNN.o1179880
```

```
9824/10000 [=====>.] - ETA: 0s
9856/10000 [=====>.] - ETA: 0s
9888/10000 [=====>.] - ETA: 0s
9920/10000 [=====>.] - ETA: 0s
9952/10000 [=====>.] - ETA: 0s
9984/10000 [=====>.] - ETA: 0s
10000/10000 [=====>.] - 41s 4ms/step
```

```
Test score: 0.0483192791523
```

```
Test accuracy: 0.99
```

```
Working directory was /user/scratch/gent/vsc425/vsc42544/KERAS_MNIST_CNN_1179880.master19.golett.gent.vsc
```

# Advanced Application Examples & Opportunities




## Deep Learning and Unsupervised Clustering for Analysis of Cellular Structures in the Human Brain

C. Bodenstern\*, H. Spitzer\*\*, P. Glock\*\*, M. Riedel\*, T. Dickscheid\*

\* High Productivity Data Processing, Jülich Supercomputing Center (JSC)  
 \*\* Big Data Analytics, Institute of Neuroscience and Medicine (INM-1)


### Cytoarchitectonic Mapping




Layer structure differs between cytoarchitectonic areas [8]. Classical methods to locate borders include image segmentation, mathematical morphology, and correlation of local intensity profiles.

#### Goals

- Investigate the potential of modern machine learning techniques to support the analysis
- Increase degree of automatization (towards high throughput processing)
- Find qualitative and quantitative measures for cellular distributions




### Gray/white matter segmentation




- Use Deep-CNNs on to predict cerebral cortex tissue
  - Masks partially available
  - Post-processing to find optimal results (e.g. Graph Cut)
  - Prediction for high resolution images parallel on HPC
- Next steps**
- Increase classifier size / Improve post-processing
  - Enhance masks

### Cell detection




- Compute shape attributes for each connected region
  - Random Forest / SVM for prediction
- Next steps**
- Split cells using clustering, parallel for each label
  - Create a unified workflow

### Parcellation of cytoarchitectonic cortical regions



- Cortical areas show different cell densities
- Using density based clustering (DBSCAN) to find regions of different density
- Scalable and parallel HPC-DBSCAN implementation to cluster large number of cells
- Compare to results with different clustering techniques

### Density clustering to reflect laminar layers



- Density based clustering of cell borders in the cortex


Enabling new Start-ups



## Automated Soccer Scene Tracking Using Deep Neural Networks

C. Bodenstern\*, M.Goetz\*\*, M. Riedel\*

\* High Productivity Data Processing, Jülich Supercomputing Center (JSC)



### Construction of an automated pipeline for the broadcast of football games

- Most matches will never be recorded e.g. amateurs
- TV camera systems and cinematographer are expensive
- Simple object tracking—i.e. the ball—is not sufficient for specific game situations like e.g. corners
- Learn the scene tracking using Deep Neural Networks
- Goal: determine the point of interest coordinates and camera zoom for each frame




Figure: Parameters of the entire field (left), the images represent the tracked areas by the operators, (middle) captured ground truth, (right) Prediction made by Deep Neural Networks

### Data Source

- Capture the entire field with multiple static cameras—two privileges with 2 or 6 images respectively
- Split images to singular panorama
- Labeled focus in x, y, point and zoom
- Currently 30 labeled football games
- ~45 TB MPEG2 compressed
- ~1500,000 Frames
- High resolution images in 2017
- More at [www.soccernetz.de](http://www.soccernetz.de)

### Why Deep Learning?

- Convolutional Neural Networks (CNNs) are state-of-the-art in image classification and object tracking [1]
- Layers abstract different things: the audience, players, the ball, etc.
- Recurrent Neural Networks retain time information from sequentially analyzed frames [2]
- Popularity realized in highly optimized tools that use GPUs

### The Learning Outcome

- Input: DNN gets a sequence of frames
- Output: Three output neurons describing x and y position of the camera focus plus the additional zoom
- Prediction close to the ground truth and appears natural
- A look into the convolutions can show learned features



### Genetic Optimization of the Deep Neural Networks



- Genetic optimization of Network Architecture on Apache Spark Cluster
- 20 different CNN 'individuals' in parallel
- Selector after three hours of learning

Figure: Fitness over generations (blue is better)

### Future Work

- Training on new, high-quality data
- Evaluation of the usage of RNNs
- Reduction of the Network's complexity to allow real-time performance—3 FPS is sufficient
- Smooth the camera motions in a post processing step

### References


[1] S. Ioffe and A. Szegedy, "Batch Normalization: Accelerating Deep Network Training by Reducing Internal Covariate Shift," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2015, pp. 448–456.  
 [2] S. Hochreiter and J. Schmidhuber, "Long Short-Term Memory," *Artificial Intelligence Magazine*, vol. 12, no. 4, pp. 1–24, 1997.  
 [3] M. Goetz, C. Bodenstern, M. Riedel, and T. Dickscheid, "Automated Soccer Scene Tracking Using Deep Neural Networks," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017, pp. 1–8.

# CNN – Neuroscience Application

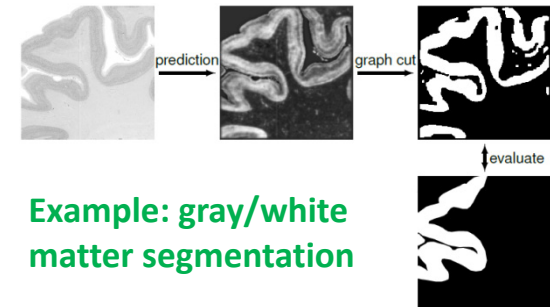


- Goal: Cytoarchitectonic Mapping

- Layer structure differs between cytoarchitectonic areas
  - Classical methods to locate borders consists of much manual work: e.g. image segmentation, mathematical morphology, etc.

- Deep Learning: Automate the process of learning ‘border features’ by providing large quantities of labelled image data
  - However: the structure setup of the deep learning network still requires manual setup (e.g. how many hidden layers, etc.)




Example: Parcellation of cytoarchitectonic cortical regions



Example: gray/white matter segmentation

Use Convolution Neural Networks: arbitrary dimension, move ‘filter’ kernel over input space, take local space into account, much cheaper, less parameters than fully connected (e.g. ANNs)


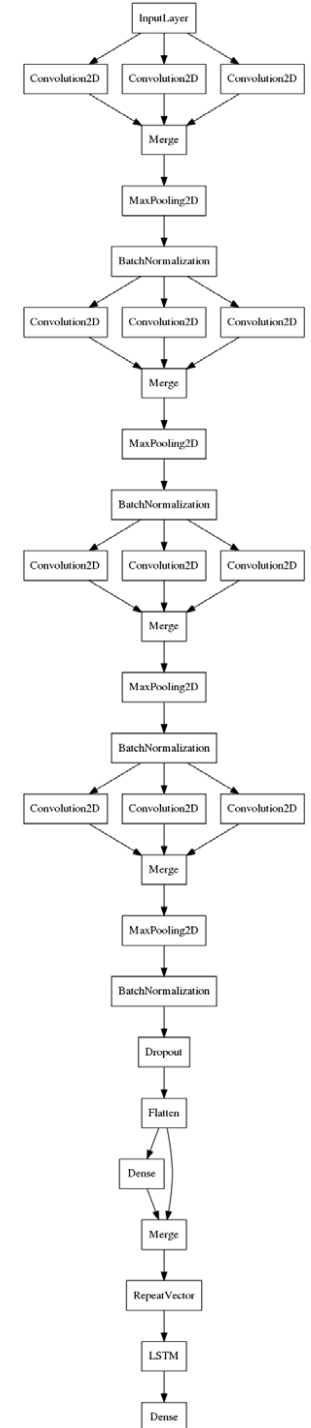



# CNN – Soccerwatch.tv Application


- Goal: Automatic zoom w/o camera man



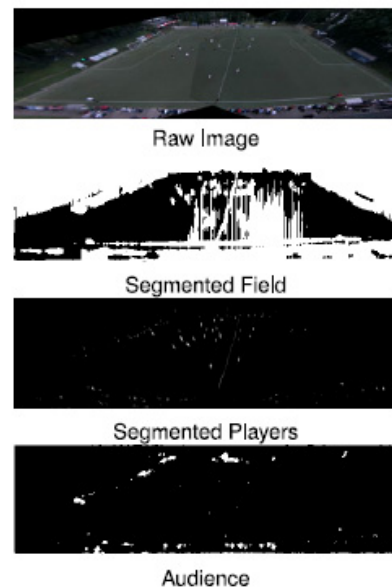
- Besides upper leagues:
  - 80k matches/week
- Recording too expensive (amateurs)
- Camera man needed
- Soccerwatch.tv provides panorama
- Approach: Find X,Y center and zoom on panorama using Deep Learning



# CNN – Soccerwatch.tv Application – Results



(Look into convolutions shows learned features)




[11] Soccerwatch.tv

# [Video] CNN Application in Autonomous Driving



*[14] YouTube Video, Speed Sign Recognition*


# Deep Learning Toolset





# Time Line of Deep Learning and Machine Learning

- Selected Frameworks only



[8] A Tour of Tensorflow



# Increasing number of Deep Learning Frameworks

- TensorFlow

- An open-source software library often used
- Supported device types are CPU and GPU

[4] *Tensorflow*



- Caffe

- Deep learning framework made with speed and modularity in mind
- Switch between CPU and GPU by setting a single flag
- E.g. train on a GPU machine, then deploy to commodity clusters

[5] *Caffe*

- Theano

- Python library for deep learning with integration of NumPY
- Transparent use of GPGPUs

[6] *Theano*

- There are a wide variety of deep learning frameworks available that support convolutional neural networks and take advantage of GPGPUs, e.g. TensorFlow, Caffe, Theano


[7] *Deep Learning Framework Comparison*

# What is a Tensor?


- Meaning
  - Multi-dimensional array used in big data analysis often today
  - Best understood when comparing it with vectors or matrices



(one dimensional tensor)  
(vector of dimension [5])



(two dimensional tensor)  
(matrix of dimensions [5,6])

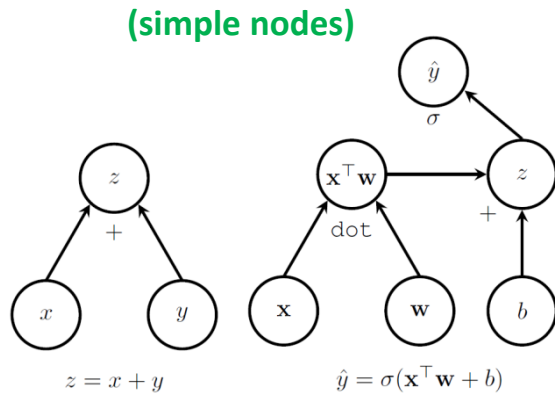


(three dimensional tensor)  
(tensor of dimension [4,4,3])

**[10] Big Data Tips, What is a Tensor?**

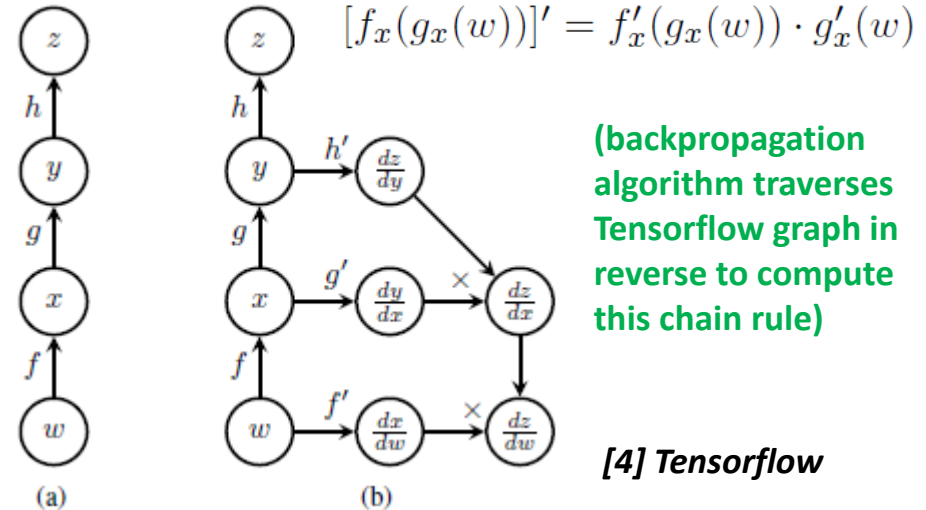
# Tensorflow Computational Graph

- **Keras as a High-Level Framework** (on top of Tensorflow)
  - Abstracts from the computational graph and **focus on layers**
- Machine learning algorithms as **computational graph**
  - Sometimes also called 'dataflow graph' to emphasize data elements
  - **Edges represent data** (i.e. often tensors) flowing between nodes
  - **Vertices / nodes are operations** of various types (i.e. combination or transformation of data flowing through the graph)



[8] A Tour of Tensorflow

(adds gradient node for each operation that takes the gradient of the previous link – outer functions – and multiplies with its own gradient)



(backpropagation algorithm traverses Tensorflow graph in reverse to compute this chain rule)

[4] Tensorflow

# Exercises – MNIST Dataset – CNN Model Check



# [Video] Backpropagation in Deep Learning Frameworks

$f(x, y) = x(y + x)$

$$\frac{\partial x(y+x)}{\partial x} = \frac{\partial x}{\partial x} \cdot (y+x) + x \cdot \frac{\partial (y+x)}{\partial x}$$


$\frac{\partial x}{\partial x} = 1$     $\frac{\partial (y+x)}{\partial x} = \frac{\partial y}{\partial x} + \frac{\partial x}{\partial x}$

$\frac{\partial y}{\partial x} = 0$     $\frac{\partial x}{\partial x} = 1$

1:09 / 4:42

[13] YouTube Video, Chain Rule in Backpropogation

# Lecture Bibliography



# Lecture Bibliography (1)

- [1] M. Nielsen, 'Neural Networks and Deep Learning',  
Online: <http://neuralnetworksanddeeplearning.com/>
- [2] Ugent Tier-2 Clusters,  
Online: <https://www.vscentrum.be/infrastructure/hardware/hardware-ugent>
- [3] A. Rosebrock, 'Get off the deep learning bandwagon and get some perspective', Online:  
<http://www.pyimagesearch.com/2014/06/09/get-deep-learning-bandwagon-get-perspective/>
- [4] Tensorflow,  
Online: <https://www.tensorflow.org/>
- [5] Cafe Deep Learning Framework,  
Online: <http://caffe.berkeleyvision.org/>
- [6] Theano Deep Learning Framework,  
Online: <http://deeplearning.net/software/theano/>
- [7] Deep Learning Framework Comparisons,  
Online: <http://neuralnetworksanddeeplearning.com/chap6.html>
- [8] A Tour of Tensorflow,  
Online: <https://arxiv.org/pdf/1610.01178.pdf>
- [9] A. Gulli and S. Pal, 'Deep Learning with Keras' Book, ISBN-13 9781787128422, 318 pages,  
Online: <https://www.packtpub.com/big-data-and-business-intelligence/deep-learning-keras>
- [10] Big Data Tips, 'What is a Tensor?',  
Online: <http://www.big-data.tips/what-is-a-tensor>
- [11] Soccerwatch.tv,  
Online: <http://www.soccerwatch.tv>



# Lecture Bibliography (2)

- [12] D. Kingma and Jimmy Ba, 'Adam: A Method for Stochastic Optimization',  
Online: <https://arxiv.org/abs/1412.6980>
- [13] YouTube Video, 'Simple explanation of how backpropagation works in deep learning libraries',  
Online: [https://www.youtube.com/watch?v=zhKWBye\\_RgE](https://www.youtube.com/watch?v=zhKWBye_RgE)
- [14] YouTube Video, 'Speed Sign Recognition by Convolutional Neural Networks',  
Online: <https://www.youtube.com/watch?v=kkha3sPoU70>

