Machinelles Lernen
«Eine kleine Einführung»

@ZimMatthias

Matthias Zimmermann
BSI Business Systems Integration AG


«Welcher Unterschied
besteht zum Deep Blue
Schachcomputer vor 20 Jahren?»
AlphaGo
Lee Sedol

AlphaGo Hardware
Powered by TPUs

(GPUs waren gestern…)

TPU: Tensor Processing Unit

Spezialhardware für maschinelles Lernen

Fast-forward 6 Jahre
in der Zukunft …

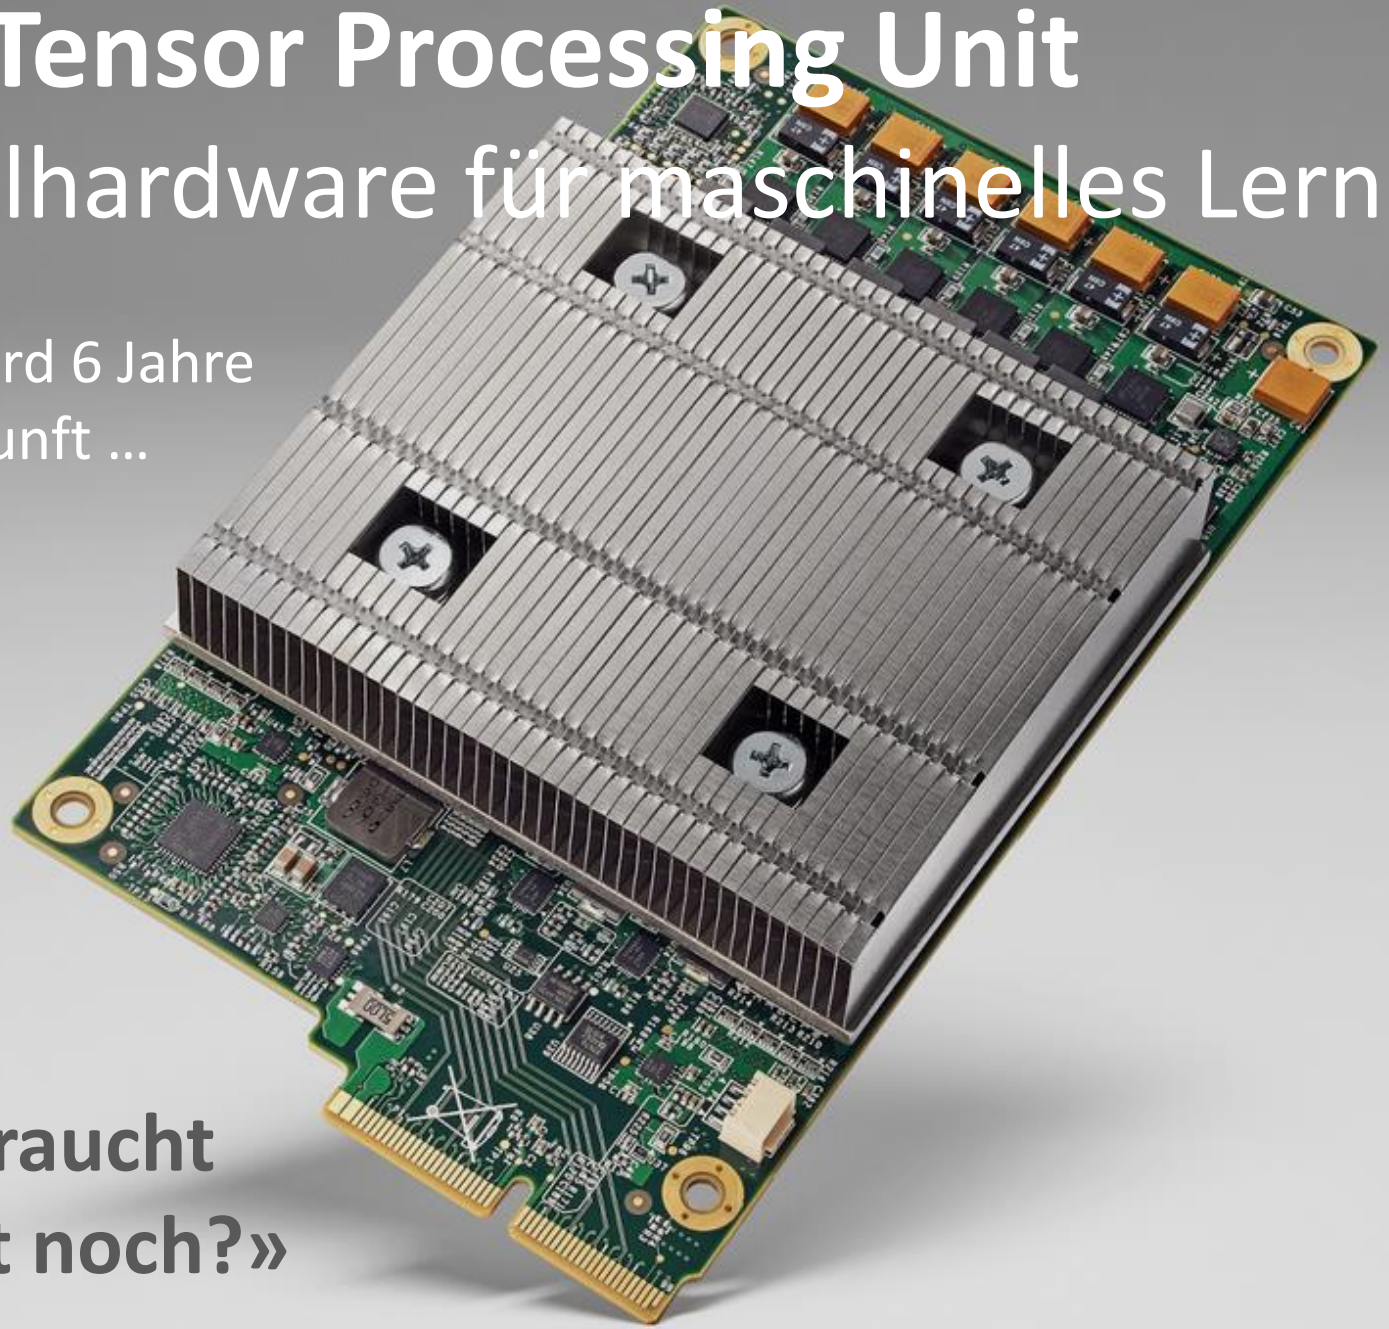

Aber …

**«Was braucht
es sonst noch?»**

nature
International weekly journal of science

2014, DeepMind

NATURE | LETTER

日本語要約

Human-level control through deep reinforcement learning

Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A. Rusu, Joel Veness, Marc G. Bellemare, Alex Graves, Martin Riedmiller, Andreas K. Fidjeland, Georg Ostrovski, Stig Petersen, Charles Beattie, Amir Sadik, Ioannis Antonoglou, Helen King, Dharshan Kumaran, Daan Wierstra, Shane Legg & Demis Hassabis

Affiliations | Contributions | Corresponding authors

Nature **518**, 529–533 (26 February 2015) | doi:10.1038/nature14236

Received 10 July 2014 | Accepted 16 January 2015 | Published online 25 February 2015

Deep Reinforcement Learning

Markov Entscheidungsprozess

➔ **Umgebung** (Atari Breakout)

➔ **Agent** der **Aktionen** ausführt (Links, Rechts, Ball starten)

➔ **Status** (Bausteien, Ort/Flugrichtung des Balls, …)

➔ **Belohnung/Rewards** (der Ball trifft einen Baustein)

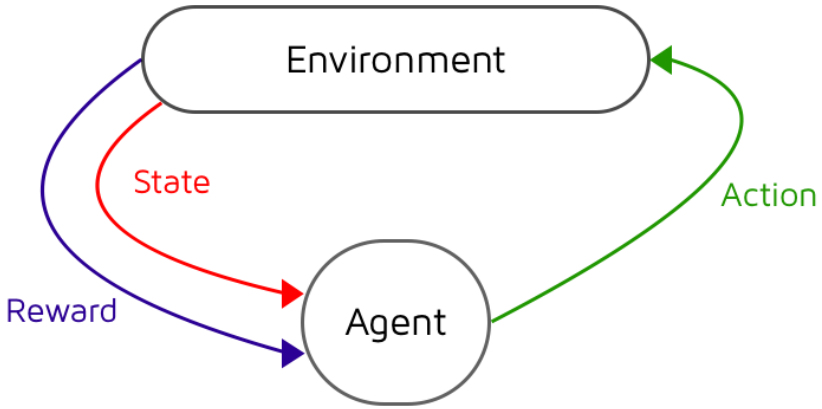

Deep Reinforcement Learning

Q-Learning (vereinfacht)

➔ Markov Entscheidungsprozess

➔ **Q(s, a)** Maximale Summe aller zukünftigen **Rewards** für **a**

```
Zufällige Initialisierung von Q
Setzen des Start-Zustands s_0
repeat
   Wähle a welche Q(s_i, a) maximiert
   Speichere neuen Zustand s_i+1 und r
   Q = update(Q, r, s_i+1)
   s_i = s_i+1
until Abbruchkriterium
```

Deep Reinforcement Learning

Deep Q-Learning (DQN)

- ➔ Q-Learning
- ➔ Q(s, a) = **Deep Neural Network** (DNN)
- ➔ Regelmässiges Nachtraining des DNN mit Erfahrungen

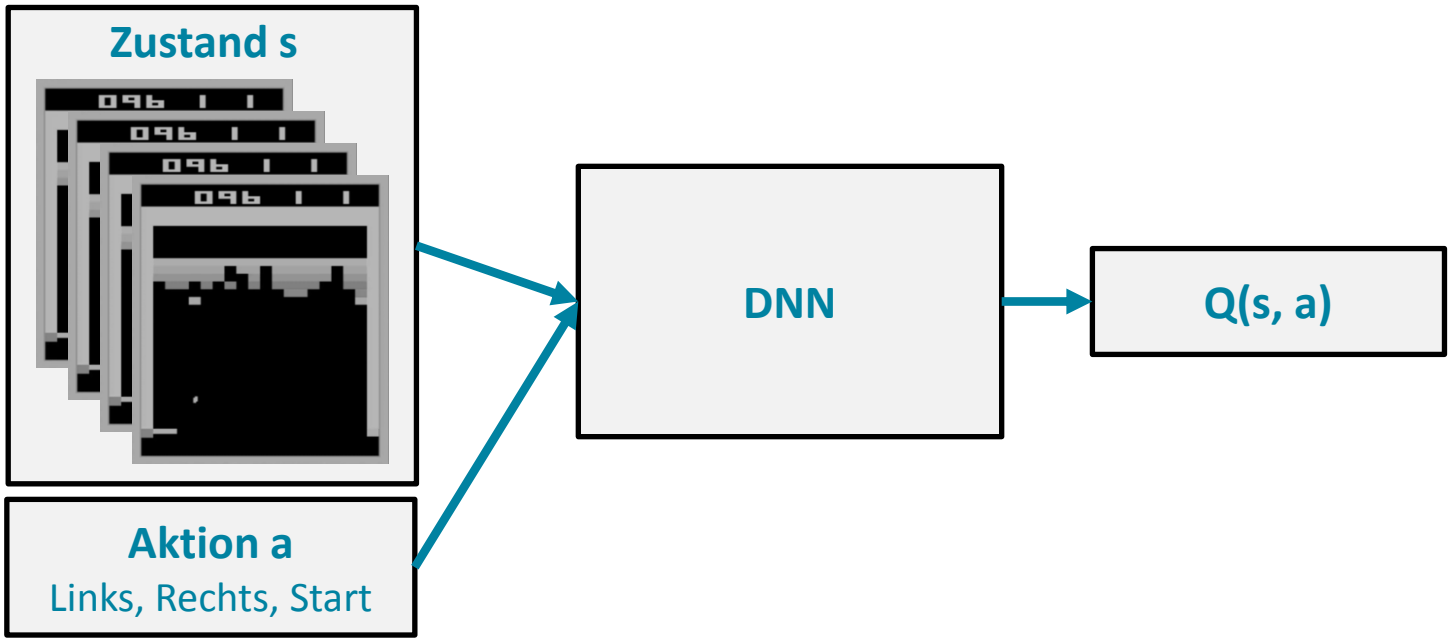

DQN
Atari Breakout

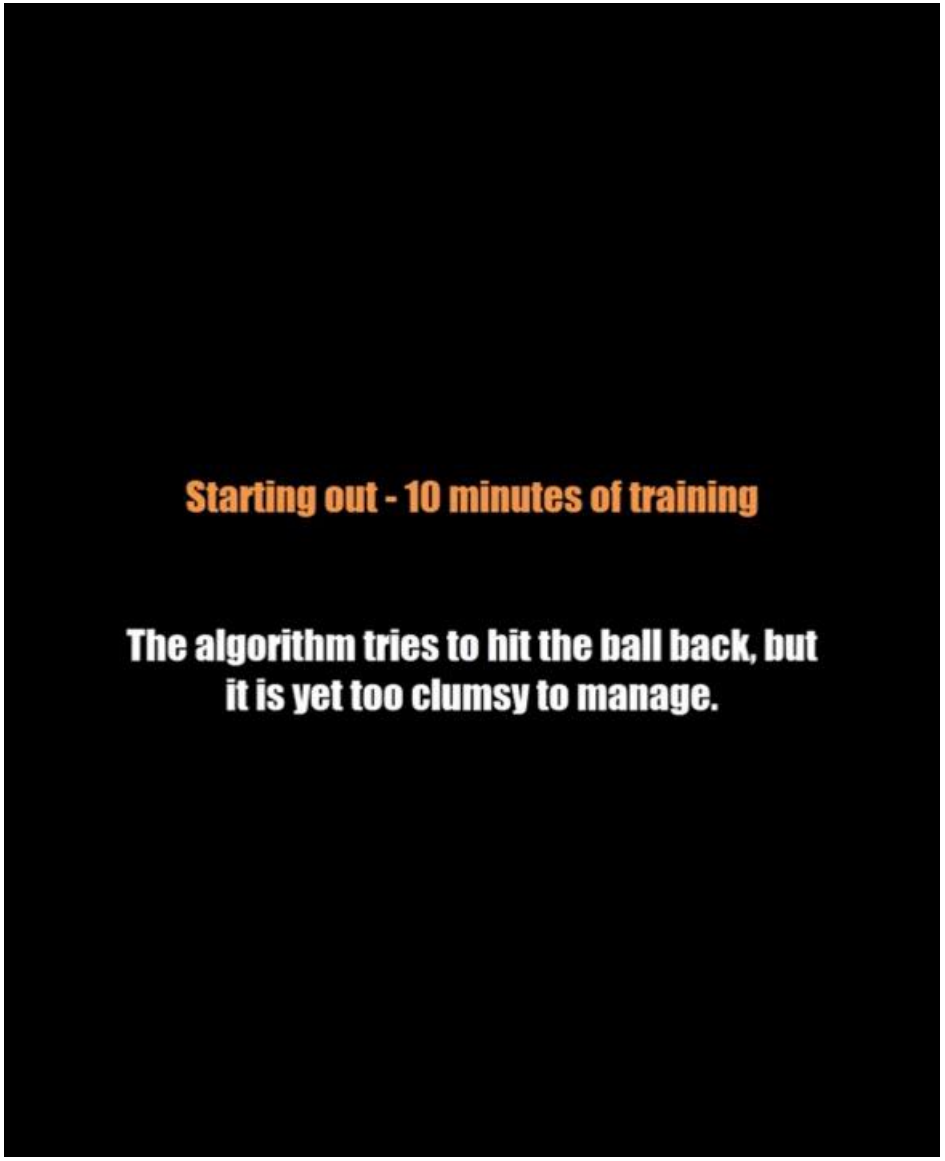

www.youtube.com/watch?v=V1eYniJ0Rnk

DeepMind and Blizzard to release StarCraft II as an AI research environment

4.11.2016

Maschinelles Lernen
Was gehört dazu?

Daten

Modelle

Training und Auswertung

ML Themen

1245 **NEGATIVE** \
shallow , noisy and pretentious .
14575 **POSITIVE** \
one of the most splendid entertainments to emerge from the french film industry in years

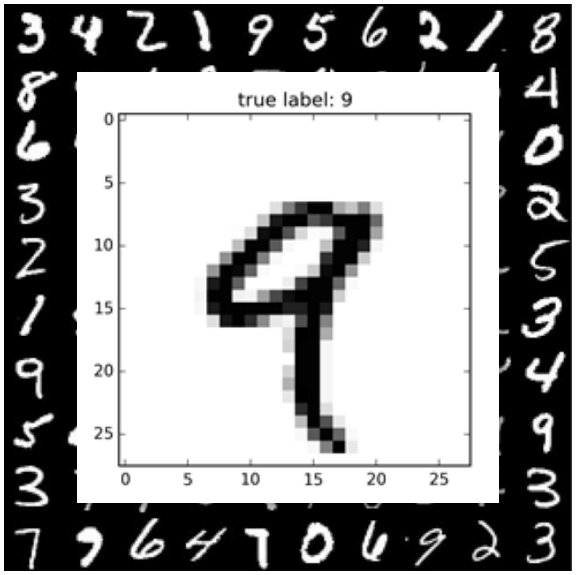

86211,**B**,12.18,17.84,77.79, …
862261,**B**,9.787,19.94,62.11, …
862485,**B**,11.6,12.84,74.34, …
862548,**M**,14.42,19.77,94.48, …
862009,**B**,13.45,18.3,86.6, …

Rosemary, Rosmarinus officinalis

Mehr zu Daten

Herausforderungen

➔ An die **RICHTIGEN** Daten ranzukommen

➔ Und **GENÜGEND** von diesen Daten zu erhalten

«Real World» Erkenntnisse

➔ Das ist nicht der spannende Teil der Projekte

➔ Beschaffung und Bearbeitung geht immer länger als geplant

➔ Daten sind ausschlaggebend

Daten

Modelle

Training und Auswertung

ML Themen

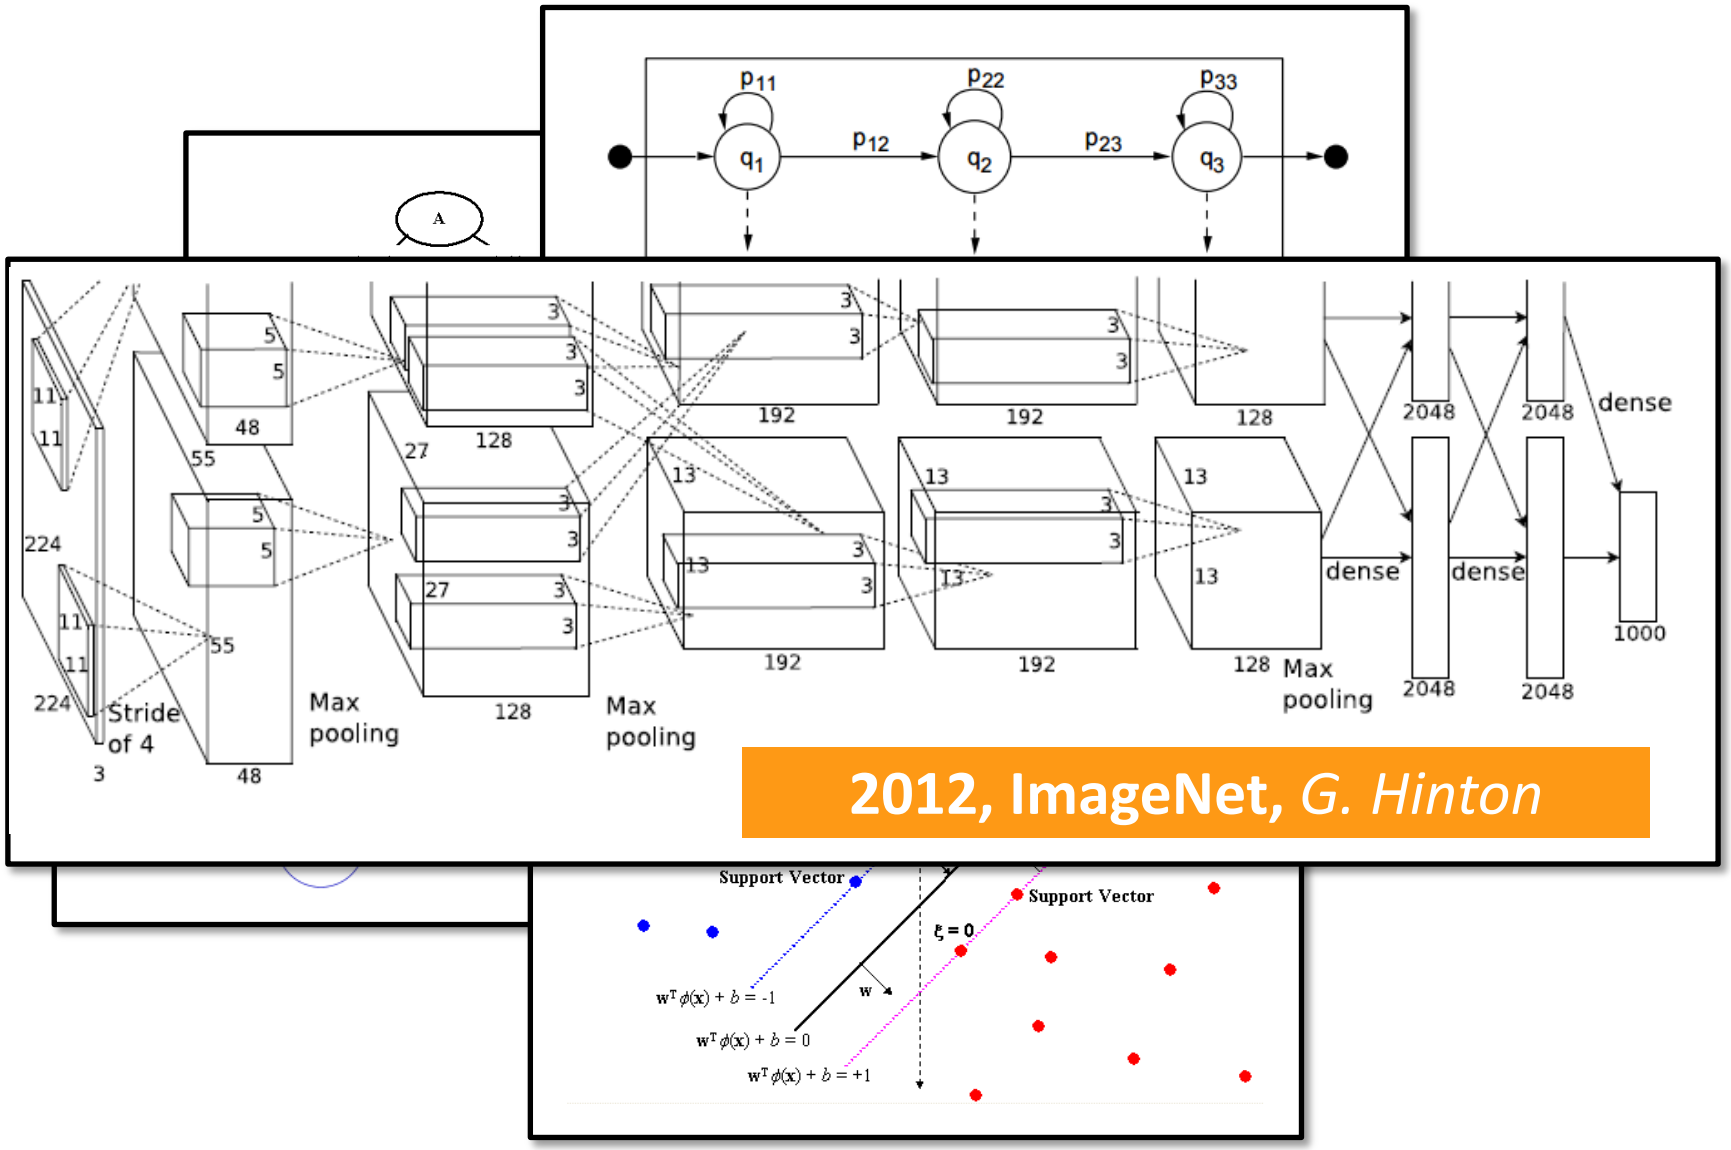
p11
p22
p33
q1
q2
q3
p12
p23
A
11
11
224
224
Stride
of 4
3
5
5
48
55
55
48
Max
pooling
27
128
3
3
3
128
Max
pooling
3
3
3
192
192
128
13
13
13
13
13
13
192
192
128
Max
pooling
dense
dense
dense
2048
2048
2048
2048
1000
2012, ImageNet, G. Hinton
Support Vector
Support Vector
ξ = 0
w
wᵀφ(x) + b = -1
wᵀφ(x) + b = 0
wᵀφ(x) + b = +1

Daten
Modelle
Training und Auswertung
ML Themen

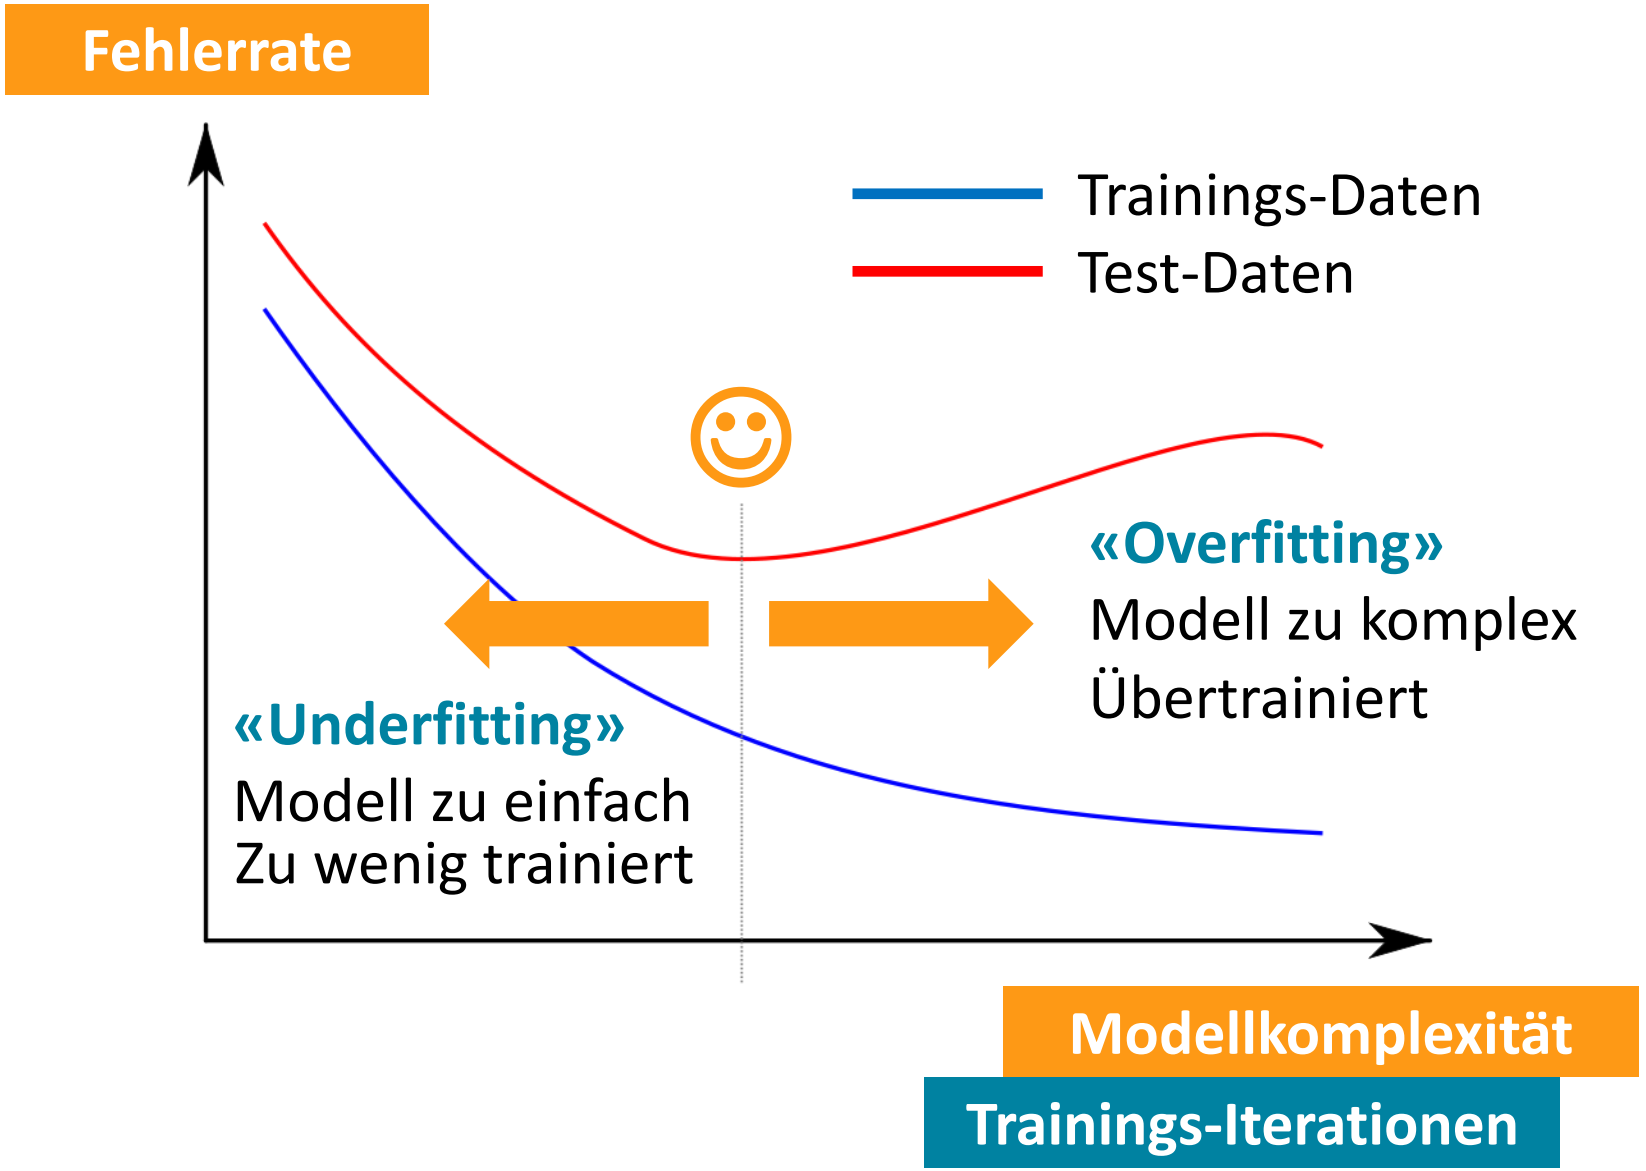
Fehlerrate
Trainings-Daten
Test-Daten
«Overfitting»
Modell zu komplex
Übertrainiert
«Underfitting»
Modell zu einfach
Zu wenig trainiert
Modellkomplexität
Trainings-Iterationen

Daten
Modelle
Training und Auswertung
ML Themen

CS 229
Machine Learning
Handout #2: Course Schedule

Syllabus

- **Introduction** (1 class)
  Basic concepts.

- Supervised learning. (7 classes)
  Supervised learning setup. LMS.
  Logistic regression. Perceptron. Exponential family.
  Generative learning algorithms. Gaussian discriminant analysis.
  Support vector machines.
  Model selection and feature selection.
  Ensemble methods: Bagging, boosting.
  Evaluating and debugging learning algorithms.

- **Learning theory.** (3 classes)
  Bias/variance tradeoff. Union and Chernoff/Hoeffding bounds.
  VC dimension. Worst case (online) learning.
  Practical advice on how to use learning algorithms.

- Unsupervised learning. (5 classes)
  Clustering. K-means.
  EM. Mixture of Gaussians.
  Factor analysis.
  PCA (Principal components analysis).
  ICA (Independent components analysis).

- Reinforcement learning and control. (4 classes)
  MDPs. Bellman equations.
  Value iteration and policy iteration.
  Linear quadratic regulation (LQR). LQG.
  Q-learning. Value function approximation.
  Policy search. Reinforce. POMDPs.

Supervised Learning
- Lernen anhand von Beispielen
- Richtige Antwort pro Beispiel bekannt

Unsupervised Learning
- Clusteranalysen
- Finden von Ausreissern in Daten

Reinforcement Learning
- Interaktion mit dynamischer Umgebung
- Belohnungsfunktion

Demos

Demos

1. Finden von Ausreissern (Brustkrebsdaten)
2. Stimmungsanalyse (Filmkritiken)
3. Mustererkennung (MNIST Bilder)

Finden von Ausreissern
Diagnostische Brustkrebsdaten

WDBC Data Set

«12.86,18,83.19,506.3,0.09934, … »

➔ Gutartig oder Bösartig?

86211,**B**,12.18,17.84,77.79, …
862261,**B**,9.787,19.94,62.11, …
862485,**B**,11.6,12.84,74.34, …
862548,**M**,14.42,19.77,94.48, …
862009,**B**,13.45,18.3,86.6, …

Modell/Algorithmus

➔ LOF (Local Outlier Factor)

Direkt Implementation in Java

➔ Java

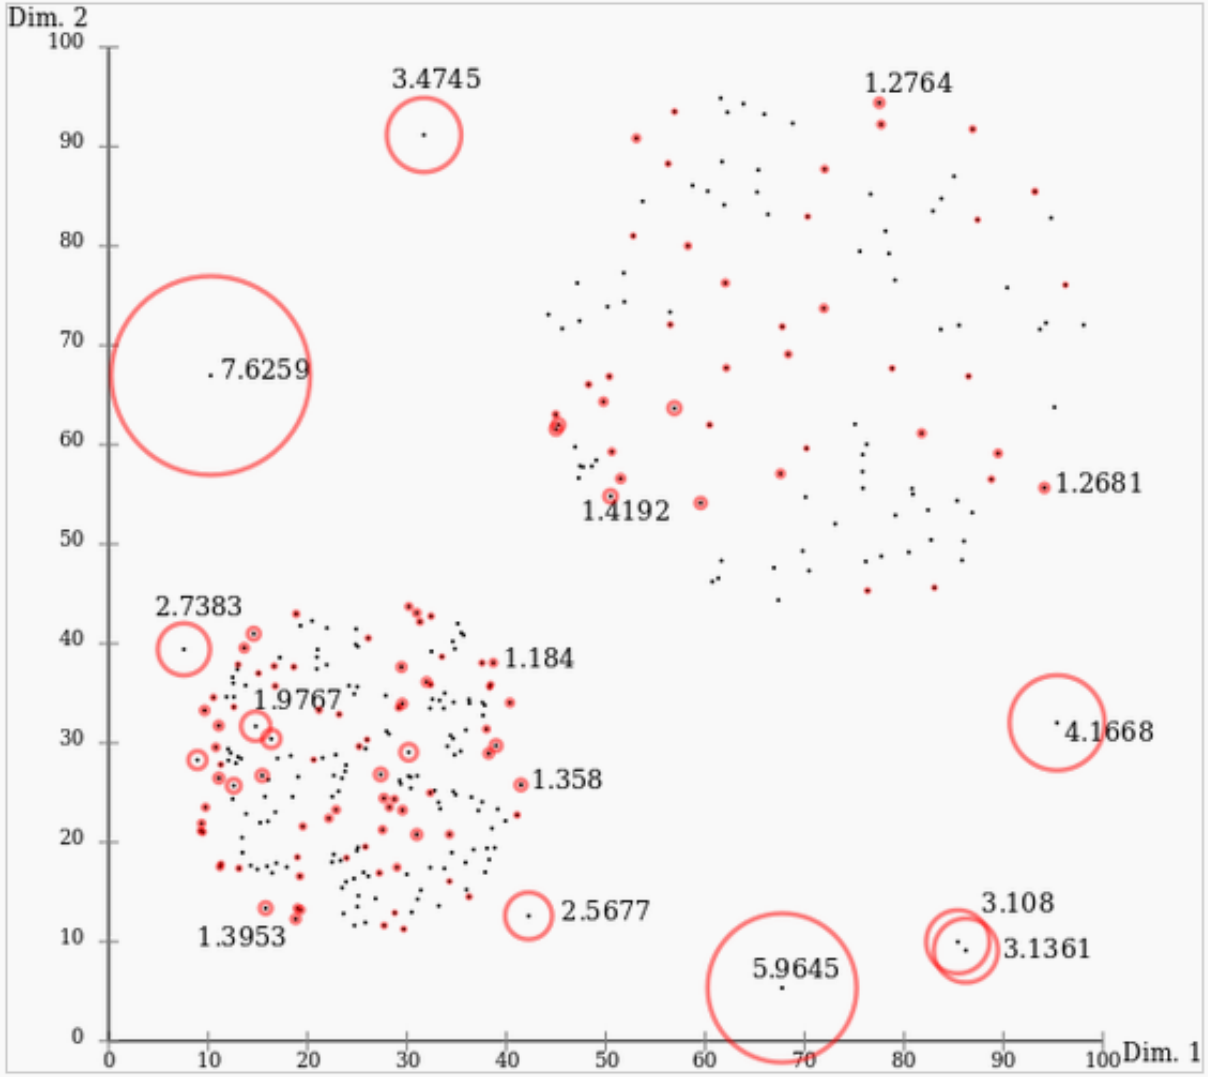

LOF-Werte visualisiert mit ELKI. Obwohl der Cluster oben rechts eine mit den Ausreißern unten links vergleichbare Dichte hat, werden sie korrekt klassifiziert.

Natural Language Processing
Stimmungsanalyse

Filmkritiken

«please, someone, stop eric schaeffer before he makes another film»

➔ Positiv oder negativ?

1245 **NEGATIVE** \
shallow , noisy and pretentious .
14575 **POSITIVE** \
one of the most splendid entertainments to emerge from the french film industry in years

Modell

➔ Naive Bayes

Mallet Toolkit

➔ Statistical NLP

➔ Open Source (CPL)

➔ Java

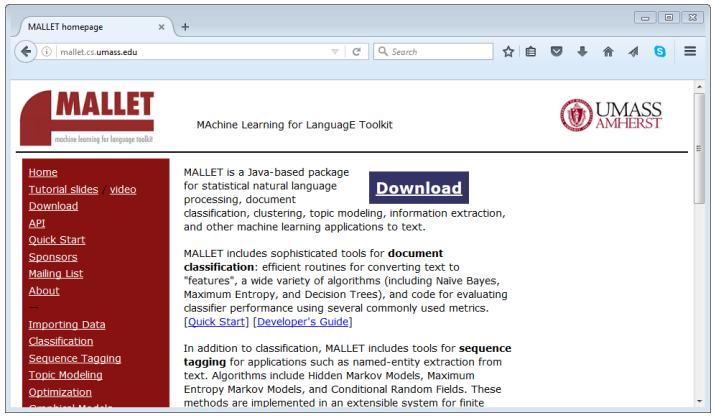

Mustererkennung
Handgeschriebene Ziffern

Daten - Handschrift

- ➔ MNIST Datenbank
- ➔ Welche Ziffer ist es?

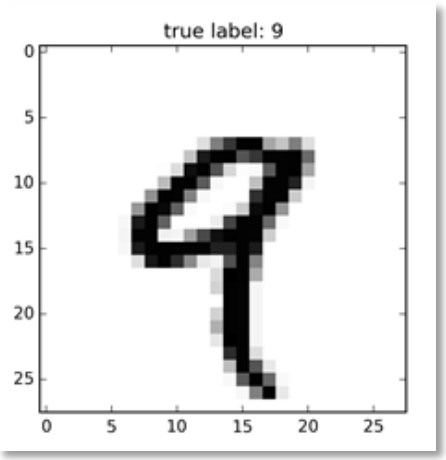

Modell

- ➔ Deep Neural Network (LeNet-5)

Deeplearning4j

- ➔ Deep Learning Library
- ➔ Open Source (Apache)
- ➔ Java

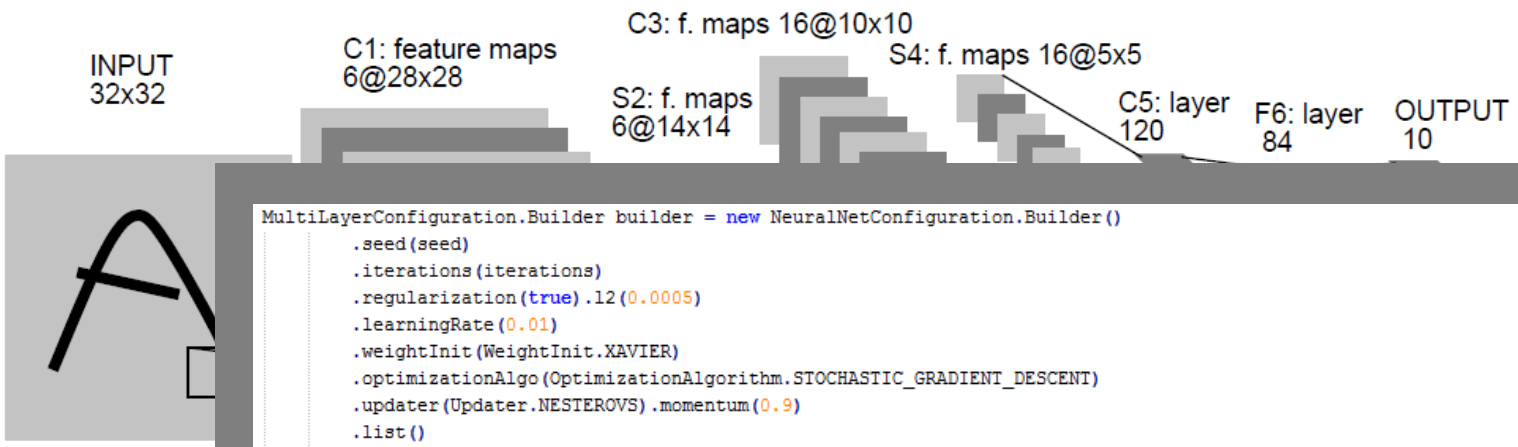


```
MultiLayerConfiguration.Builder builder = new NeuralNetConfiguration.Builder()
        .seed(seed)
        .iterations(iterations)
        .regularization(true).l2(0.0005)
        .learningRate(0.01)
        .weightInit(WeightInit.XAVIER)
        .optimizationAlgo(OptimizationAlgorithm.STOCHASTIC_GRADIENT_DESCENT)
        .updater(Updater.NESTEROVS).momentum(0.9)
        .list()
        .layer(0, new ConvolutionLayer.Builder(5, 5)
                //nIn and nOut specify depth. nIn here is the nChannels and nOut is the number of filters to be applied
                .nIn(nChannels)
                .stride(1, 1)
                .nOut(20)
                .activation("identity")
                .build())
        .layer(1, new SubsamplingLayer.Builder(SubsamplingLayer.PoolingType.MAX)
                .kernelSize(2,2)
                .stride(2,2)
                .build())
        .layer(2, new ConvolutionLayer.Builder(5, 5)
                //Note that nIn need not be specified in later layers
                .stride(1, 1)
                .nOut(50)
                .activation("identity")
                .build())
        .layer(3, new SubsamplingLayer.Builder(SubsamplingLayer.PoolingType.MAX)
                .kernelSize(2,2)
                .stride(2,2)
                .build())
        .layer(4, new DenseLayer.Builder().activation("relu")
                .nOut(500).build())
        .layer(5, new OutputLayer.Builder(LossFunctions.LossFunction.NEGATIVELOGLIKELIHOOD)
                .nOut(outputNum)
                .activation("softmax")
                .build())
        .setInputType(InputType.convolutionalFlat(28,28,1)) //See note below
        .backprop(true)
        .pretrain(false);
```

n connections

Fig. 2. Architec … e. a set of units whose weigh…

1998, … LeCun

Die letzten 4 Jahre

Die letzten 4 Jahre

2012 **ImageNet** (U Toronto)

2014 **Automatic Image Captioning** (U Stanford, …)

2015 **TensorFlow** (Google)

2016 **DNC** (Google DeepMind)

2014, Stanford

Deep Visual-Semantic Alignments for Generating Image Descriptions

Andrej Karpathy Li Fei-Fei
Department of Computer Science, Stanford University
{karpathy,feifeili}@cs.stanford.edu

Abstract

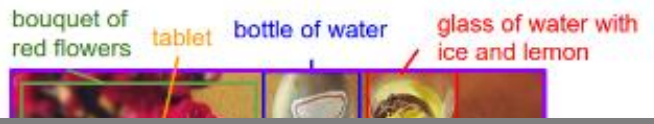

construction worker in orange safety vest is working on road.

two young girls are playing with lego toy.

boy is doing backflip on wakeboard.

TensorFlow ™

GET STARTED TUTORIALS HOW TO MOBILE API RESOURCES ABOUT

Fork me on GitHub

TensorFlow is an Open Source Software Library for Machine Intelligence

GET STARTED

2015, Google

About TensorFlow

TensorFlow™ is an open source software library for numerical computation using data flow graphs. Nodes in the graph represent mathematical operations, while the graph edges represent the multidimensional data arrays (tensors) communicated between them. The flexible architecture allows you to deploy computation to one or more CPUs or GPUs in a desktop, server, or mobile device with a single API. TensorFlow was originally developed by researchers and engineers working on the Google Brain Team within Google's Machine Intelligence research organization for the purposes of conducting machine learning and deep neural networks research, but the system is general enough to be

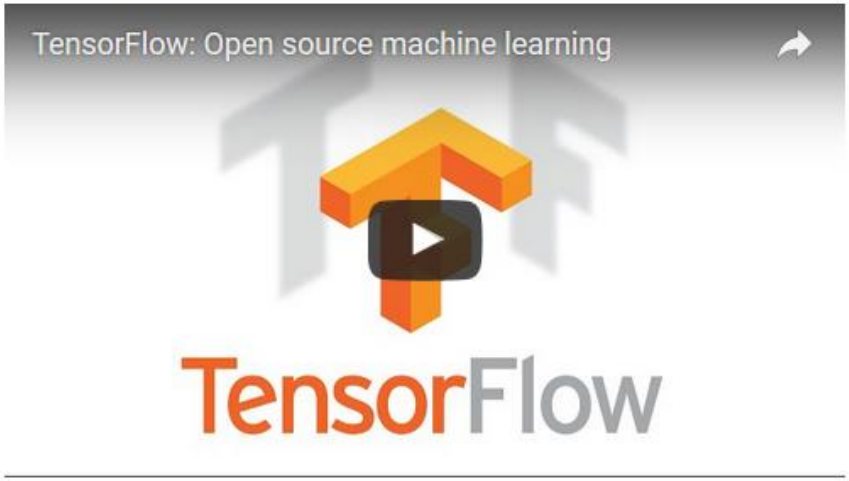

Diff

Mind

Random Training Graph

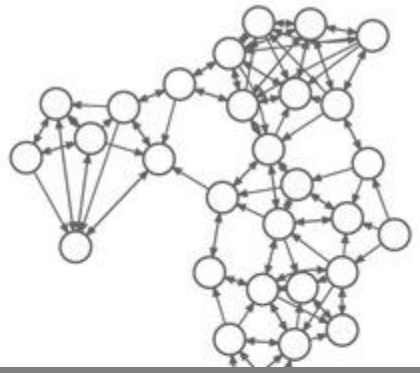

London Underground

The whole system is differentiable, and can therefore be trained end-to-end with gradient descent, allowing the network to learn how to operate and organize the memory in a goal-directed manner.

(BakerSt, OxfordCircus, Bakerloo)
...
(LeicesterSq, CharingCross, Northern)
(TottenhamCtRd, LeicesterSq, Northern)
(OxfordCircus, PiccadillyCircus, Bakerloo)
(OxfordCircus, NottingHillGate, Central)
(OxfordCircus, Euston, Victoria)

- 84 edges in total

Answer:
(BondSt, NottingHillGate, Central)
(NottingHillGate, GloucesterRd, Circle)
...
(Westminster, GreenPark, Jubilee)
(GreenPark, BondSt, Jubilee)

Answer:
(Moorgate, Bank, Northern)
(Bank, Holborn, Central)
(Holborn, LeicesterSq, Piccadilly)
(LeicesterSq, PiccadillyCircus, Piccadilly)

In a *k*
calle
answer qu *stories,*
family tree *n solve a*
block puzz

Plato liker , would
remain fix minds
can be shaped and reshaped by experience. But the wax of our memories does not just

Ausblick

Artificial Intelligence and Robotics + Add to myFT

AI and robots threaten to unleash mass unemployment, scientists warn

Intelligent machines will soon replace human workers in all sectors of economy

Read latest:

Market grows for 'regtech', or AI for regulation

OCTOBER 14, 2016

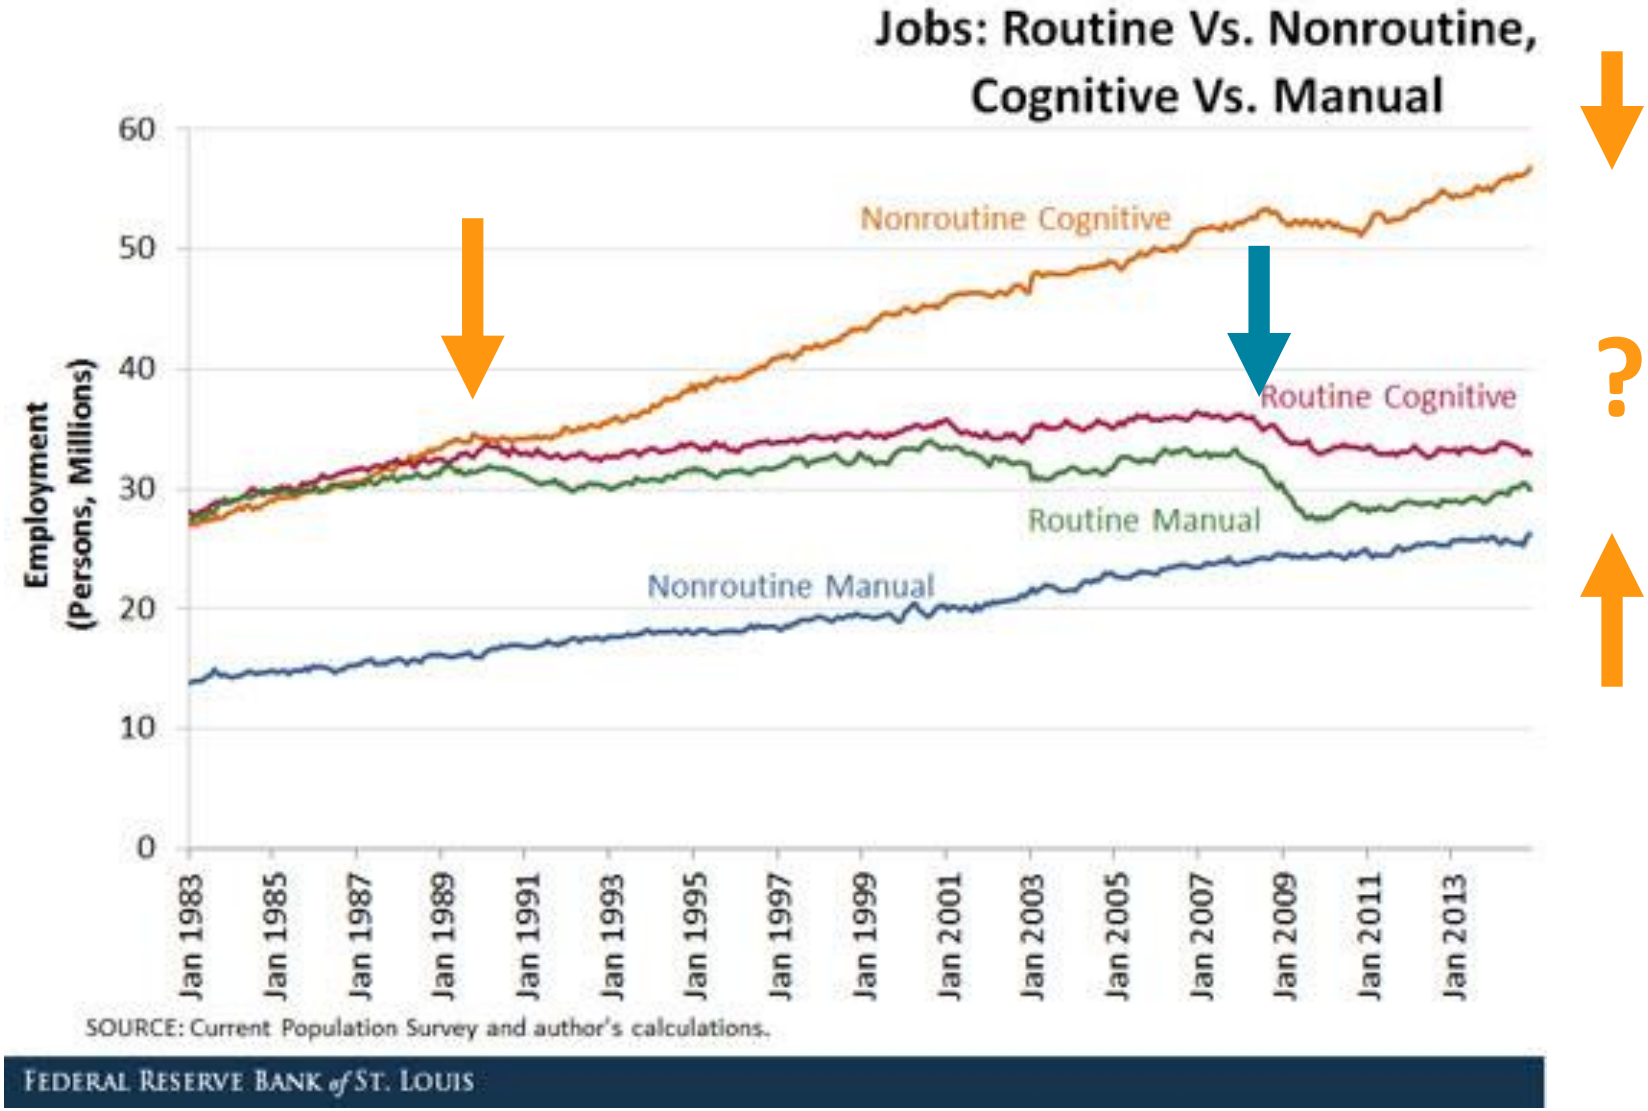
Jobs: Routine Vs. Nonroutine,
Cognitive Vs. Manual
Employment
(Persons, Millions)
60
50
40
30
20
10
0
Nonroutine Cognitive
Routine Cognitive
Routine Manual
Nonroutine Manual
Jan 1983
Jan 1985
Jan 1987
Jan 1989
Jan 1991
Jan 1993
Jan 1995
Jan 1997
Jan 1999
Jan 2001
Jan 2003
Jan 2005
Jan 2007
Jan 2009
Jan 2011
Jan 2013
?
SOURCE: Current Population Survey and author's calculations.
FEDERAL RESERVE BANK of ST. LOUIS

Eine positive Sicht

Lee Sedol

AlphaGo had opened his eyes to a new side of the game. “I have improved already,” Lee said. “It has given me new ideas.” He has not lost a match since.

Take Home Messages

Zusammenfassung

- ➔ «Deep Learning» stellt ein Durchbruch für ML / Robotik dar
- ➔ ML Systeme sind heute für viele Aufgaben praxistauglich

Was bringt die Zukunft?

- ➔ ML wird einen grossen Einfluss auf die Gesellschaft haben
- ➔ Der Übergang wird eine Herausforderung
- ➔ Zuschauen ist keine gute Option
- ➔ ML ist für alle offen

Vielen Dank!

@ZimMatthias

Matthias Zimmermann
BSI Business Systems Integration AG