Shallow Dive into Deep Learning
Eine Einführung in das „tiefe“ Lernen

Stefan H.

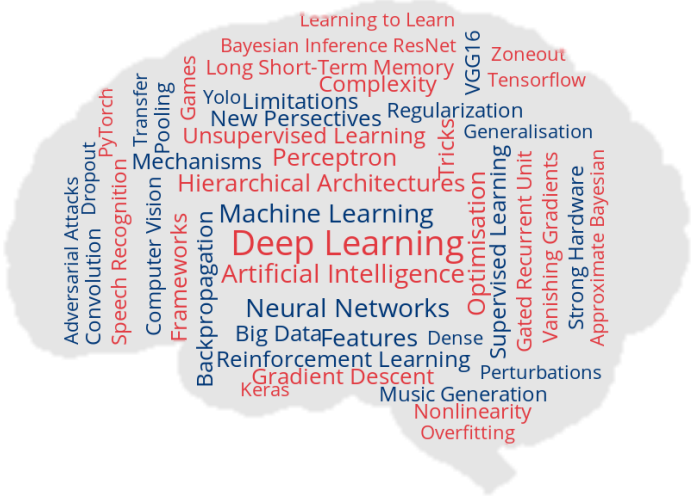

Hintergrund

Deep Learning

Künstliche Intelligenz

Computer kann menschliches Verhalten nach-bilden

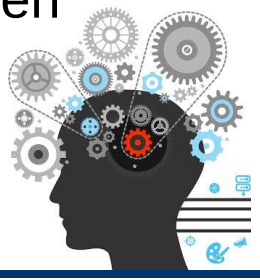

Maschinelles Lernen

Lernt ohne expliziert programmiert zu werden

Deep Learning

Lernt Merkmale in Daten mit Neuronalen Netzen

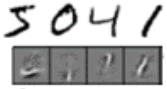

Warum reden alle darüber?

Google UK

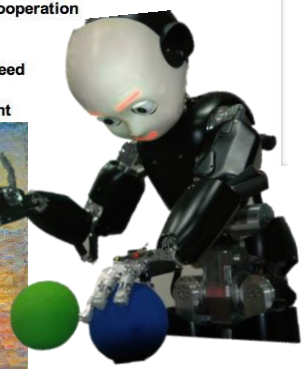

Neuronale Netze: Was ist das?

Ziel: Lernender Algorithmus

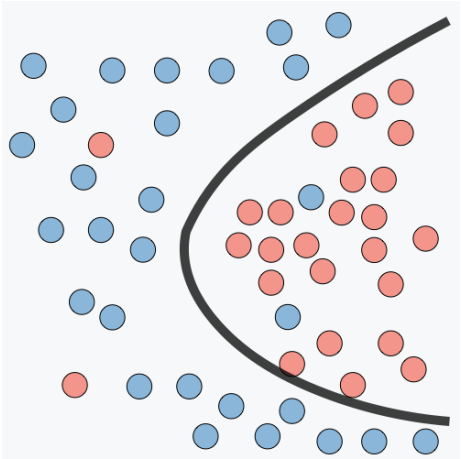

Klassifikation

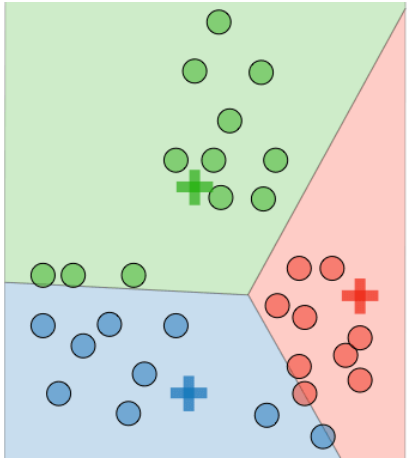

Strukturen finden

und viel mehr…

Inspiration: Gehirn

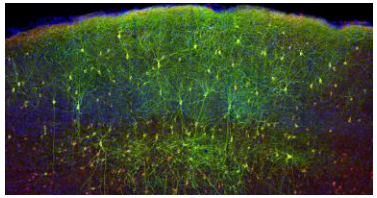

Neurone
im Kortex

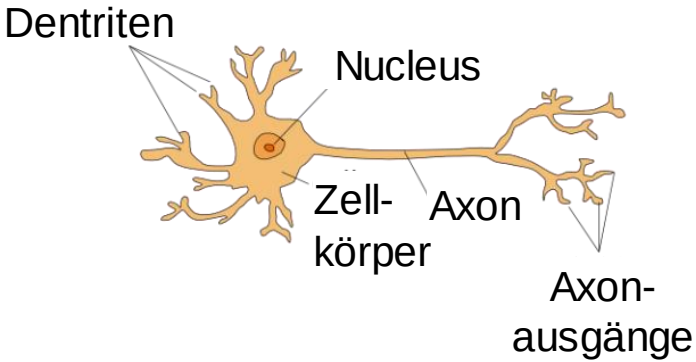

Neuronale Netze: Warum jetzt?

1952 Gradienten-abstieg

1958 Perzeptron

1986 Back-propagation

1995 Deep Convolutional NN

Big Data

- Große Datensätze
- Einfache Daten-sammlung

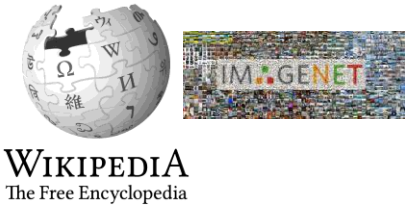

Hardware

- Grafikkarten (GPUs)
- Massive parallelsier-barkeit

Software

- Verbesserte Methoden
- Neue Modelle
- Toolboxen

Merkmale (Features)

- Von Hand entwickelte Merkmale: aufwendig, anfällig, skaliert nicht

⇨ Besser Lernen?

Low Level	Mid Level	High Level
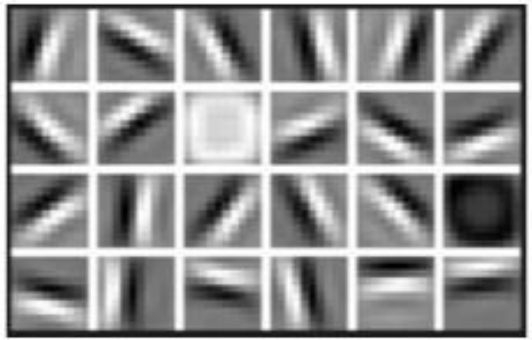		
Linien und Kanten	Augen, Nase, Ohren	Gesichter

So funktionierts: Perzeptron

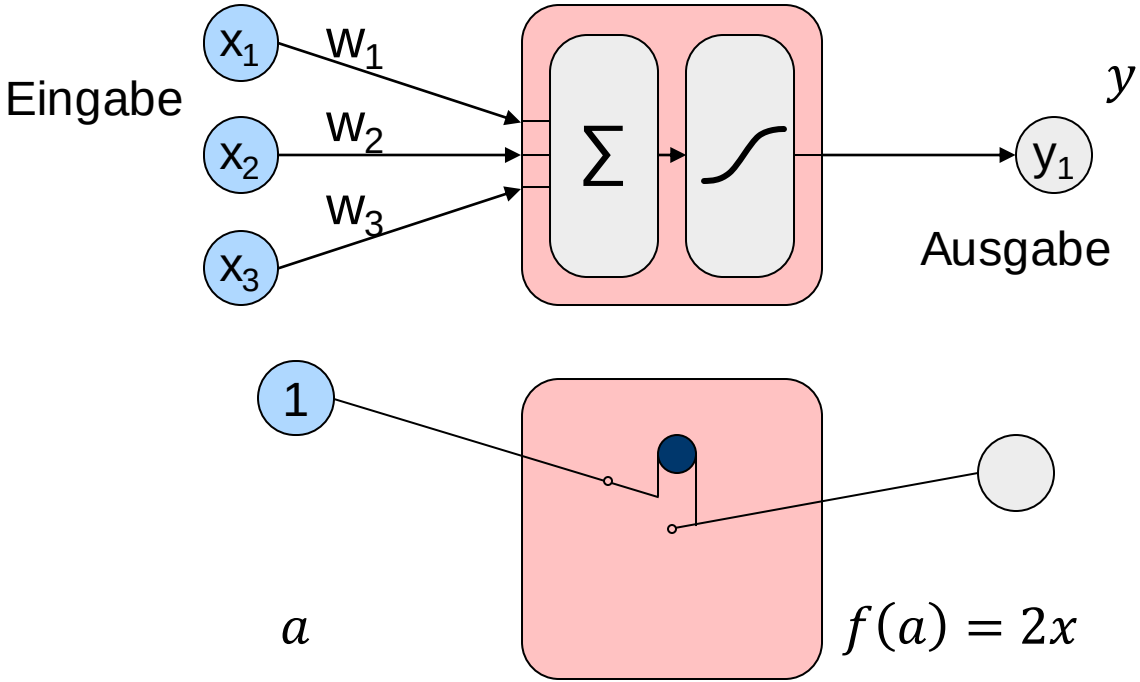

$$y = f(\mathbf{X} \cdot \mathbf{W}) = \tanh\left(\sum_{k=1}^{n} x_k w_k\right)$$

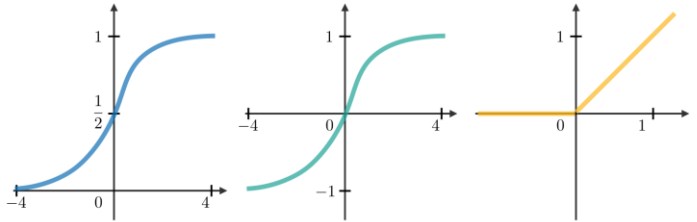

Nichlinearität

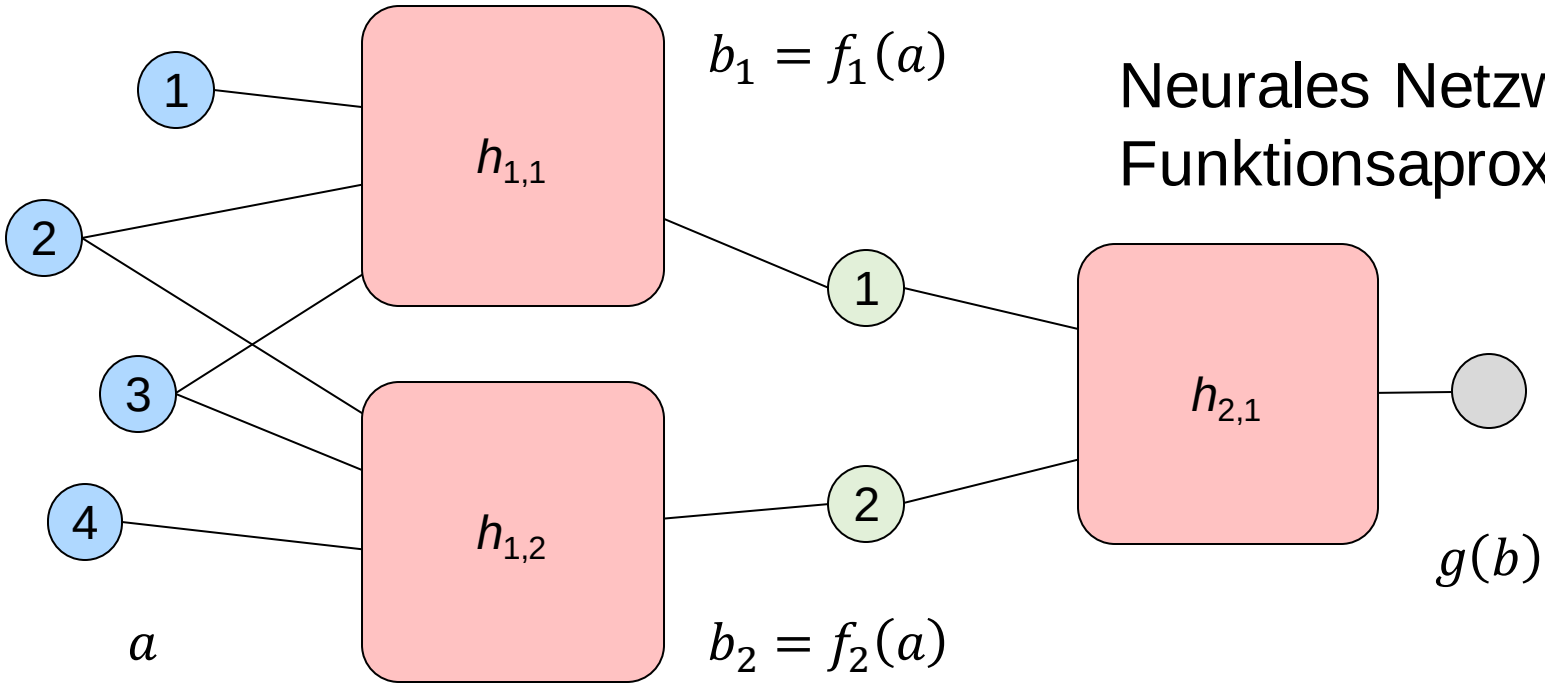

Neurales Netzwerk:
Funktionsaproxximator

Gradientenabstieg - Backprop

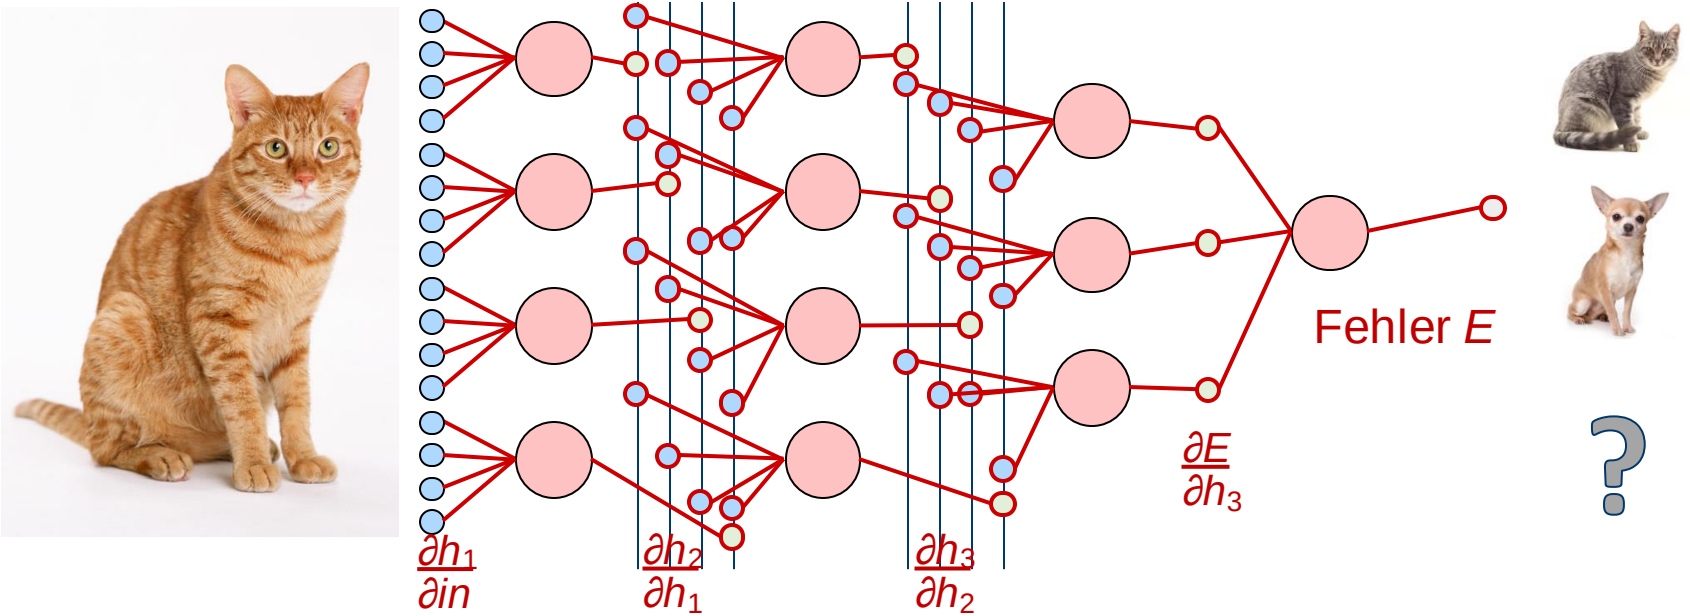

Daten-Augmentierung

- Problem: Geringe Anpassung

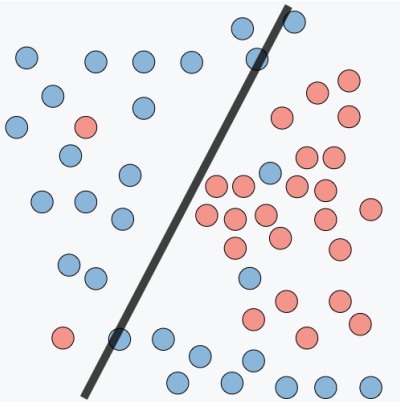

- Lösung: mehr Daten!

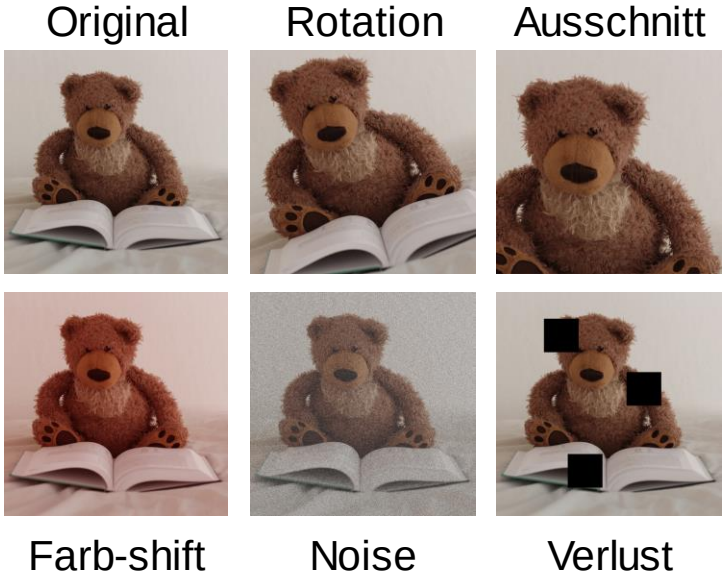

Regularisierung

- Problem: Überanpassung

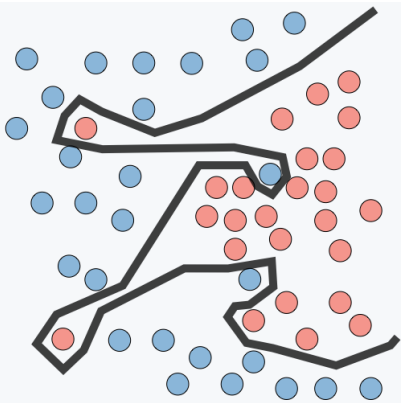

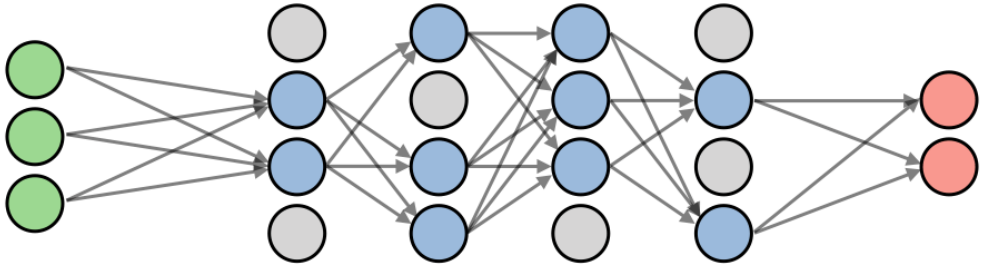

Dropout

- Lösung: Regularisierung

CNN für Bildverarbeitung

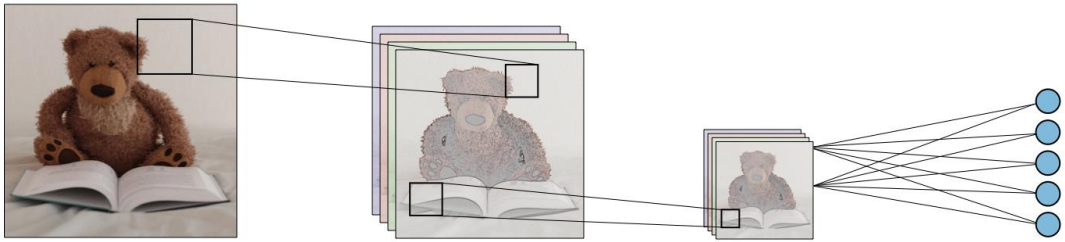

Eingabe Convolutions Pooling Dense

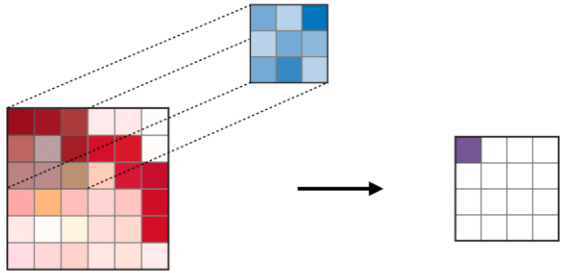

Convolution

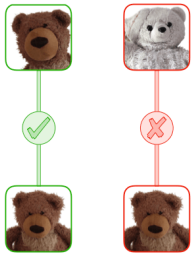

Beispiel:
Gesichter
verifizieren
o. erkennen

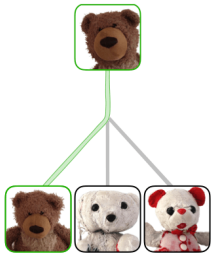

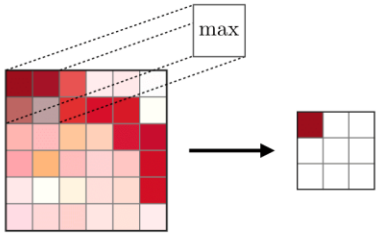

(max) Pooling

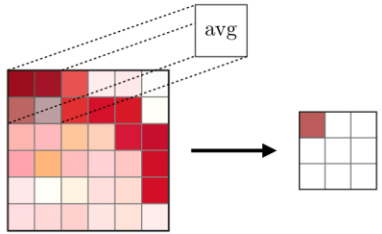

(avg) Pooling

CNN für Bildverarbeitung - Aktuell

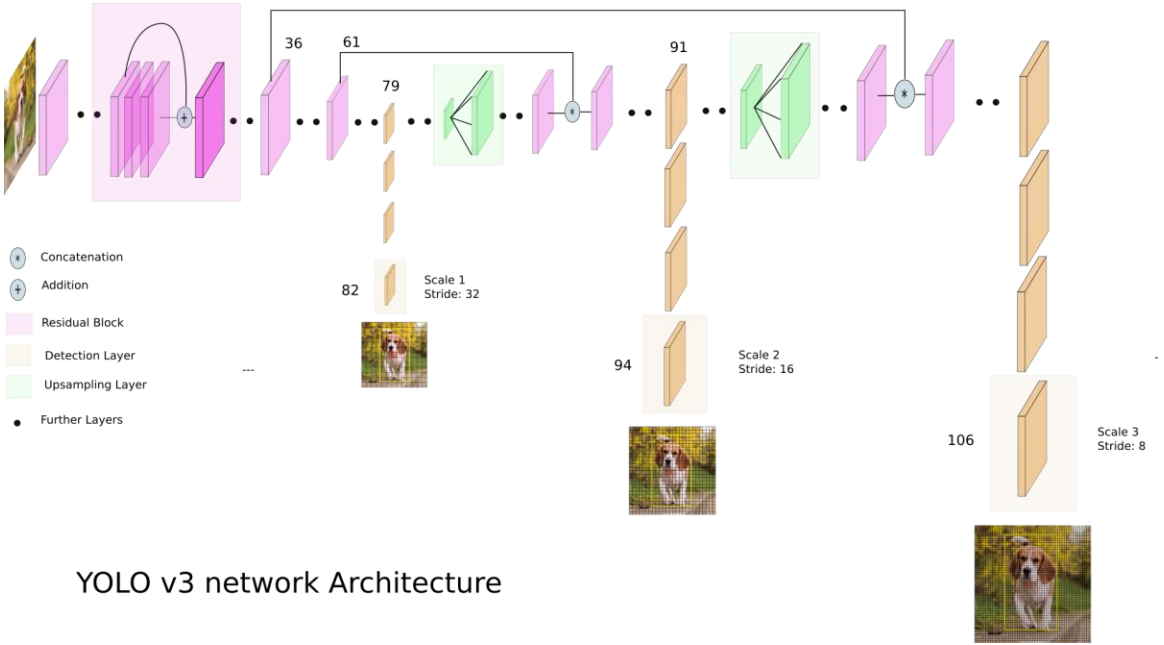

YOLO v3 network Architecture

LSTM für Sprachverarbeitung

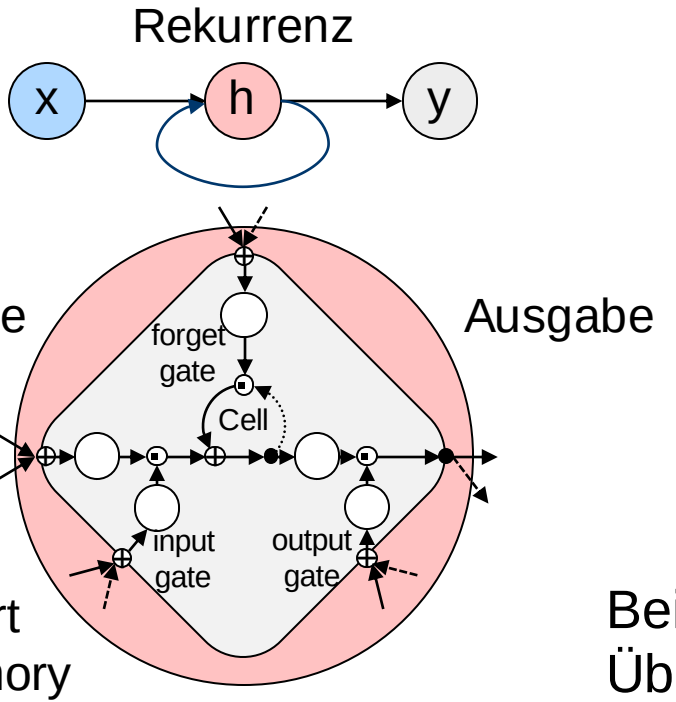

Long-Short Term Memory

Beispiel: Übersetzung

Limitierungen

Diskriminierung

Amazon Recruiting

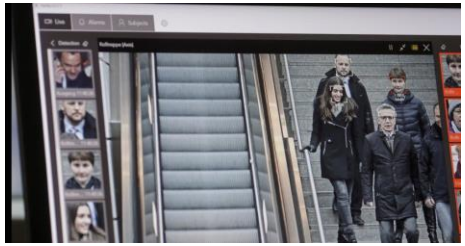

Südkreuz Überwachung

Microsoft Tay Chatbot

Adversarial attacks

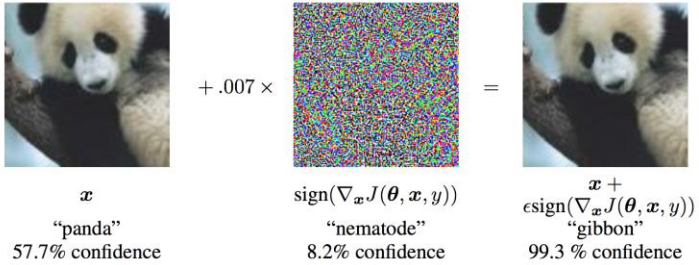

Überwachung mit Biometrie

Autonome Fahrzeuge

Perspektiven

Transfer- und Meta-Lernen

Kontext + Stil = Ergebnis

Bass Spur Generation

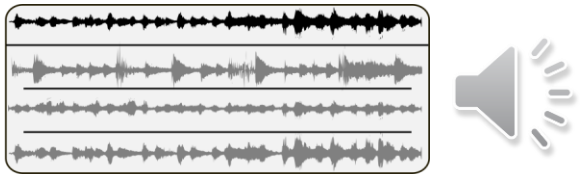

⇨ Lernen zu lernen

Auswirkungen auf Mobil- und Bio-Technologie

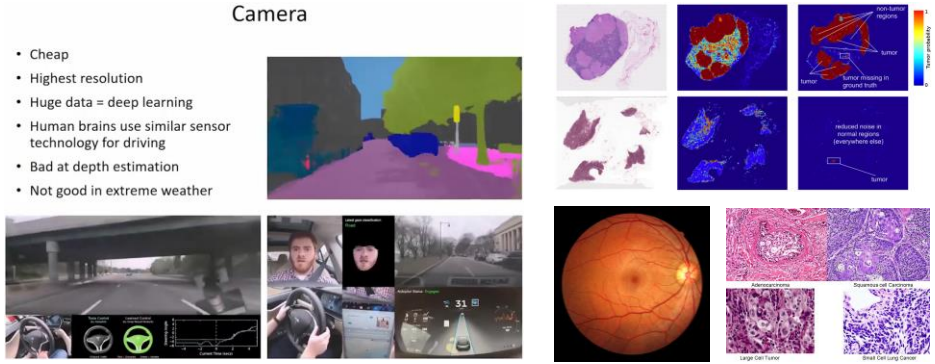

⇨ Mustererkennung anwendungsreif

Selber Loslegen

Das war’s für heute. Danke für die Aufmerksamkeit.

Lernen

- MIT 6.S191: Deep Learning (Lecture bei YouTube)
- Goodfellow. Deep Learning. MIT Press, 2017.

Ausprobieren

- https://playground.tensorflow.org
- Keras
- pyTorch
- Tensorflow