Zürcher Hochschule
für Angewandte Wissenschaften

Wie denken denkende Maschinen?
Swiss ICT Symposium, KKL Luzern, 14.11.2017

Thilo Stadelmann

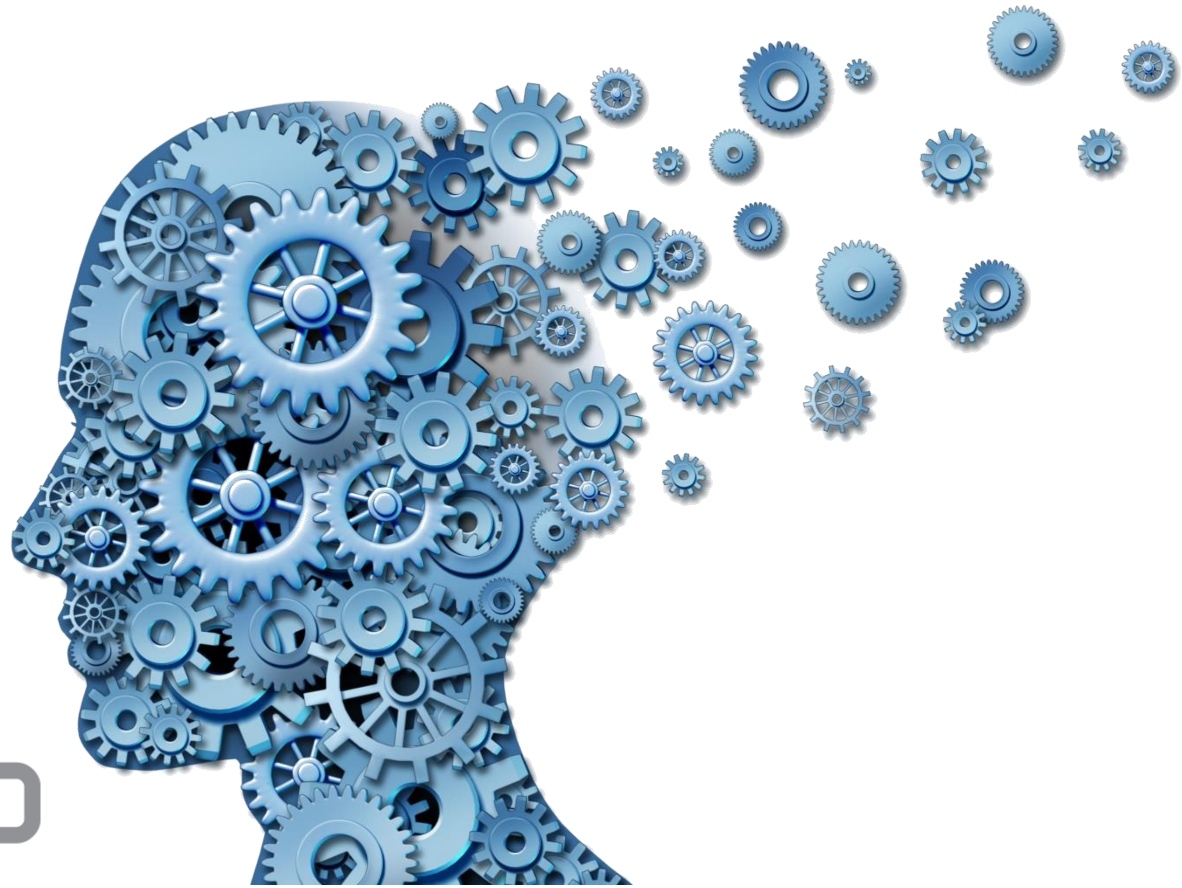

Swiss Alliance for
Data-Intensive Services

swiss group for artificial intelligence
and cognitive science

www.zhaw.ch/datalab

Was?→ Wie? → Wo?

Zürcher Hochschule
für Angewandte Wissenschaften

1

Was ist passiert?
(Eine kurze Geschichte der letzten Monate)

Google Acquires Artificial Intelligence Startup DeepMind For More Than $500M

Posted Jan 26, 2014 by Catherine Shu (@catherineshu)

Google will buy London-based artificial intelligence company DeepMind. The Information reports that the acquisition price was more than $500 million, and that Facebook was also in talks to buy the startup late last year. DeepMind confirmed the acquisition to us, but couldn't disclose deal terms.

The acquisition was originally confirmed by Google to Re/code.

Google Acquires Artificial Intelligence Startup DeepMind For More Than $500M

Posted Jan 26, 2014 by Catherine Shu (@catherineshu)

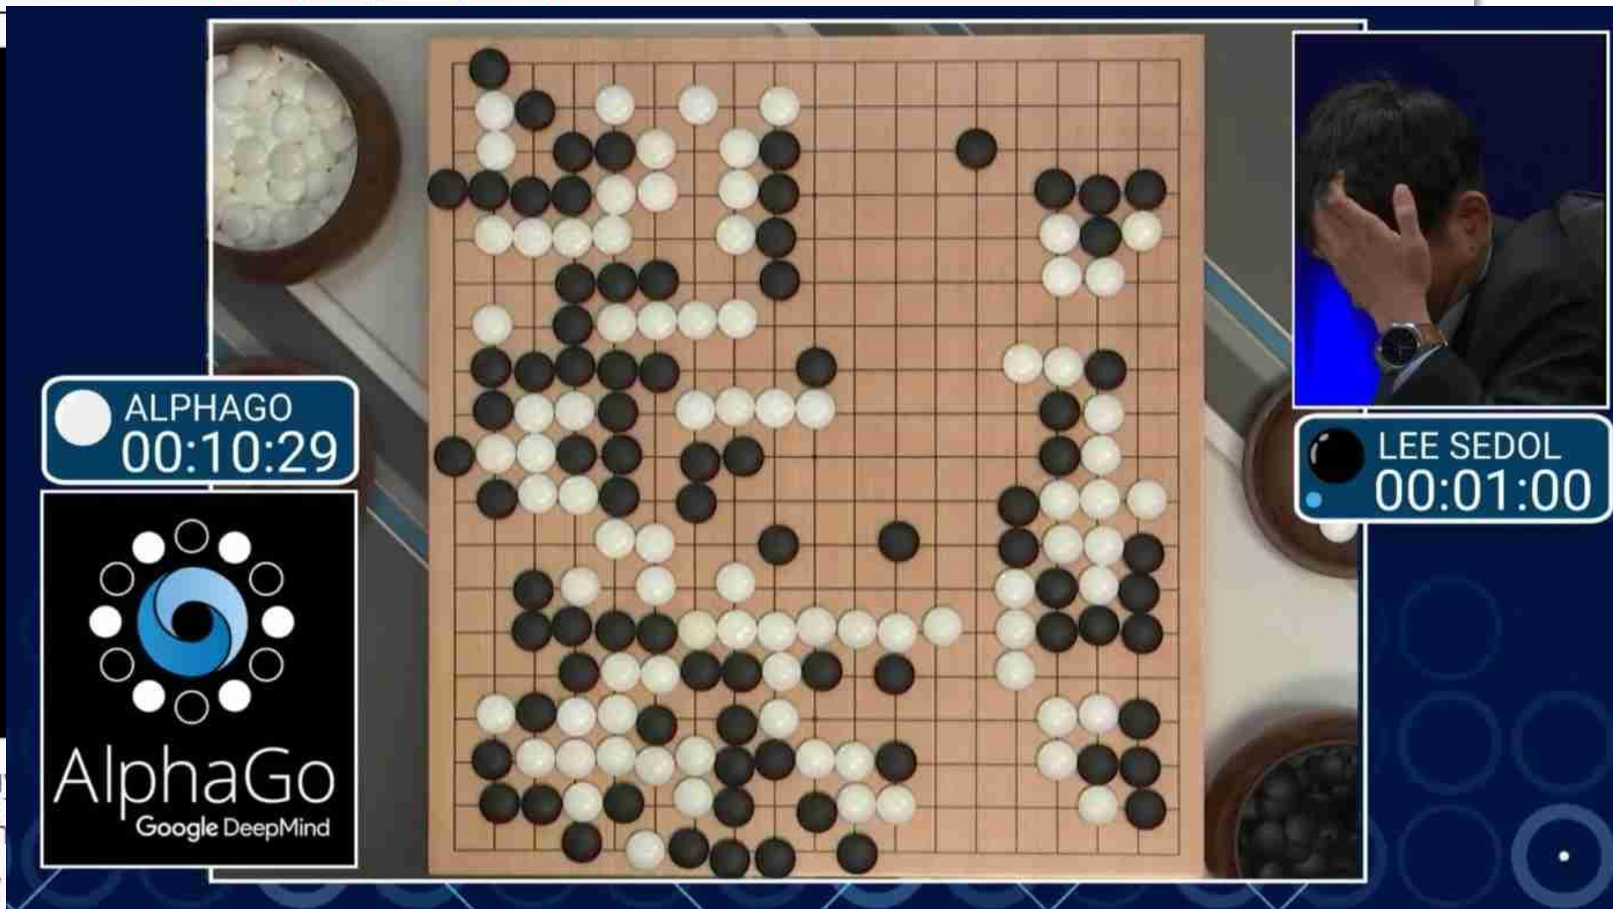

Google will bu
reports that th
in talks to buy
couldn't disclose deal terms.

The acquisition was originally confirmed by Google to Re/code.

Google Acquires Artificial Intelligence Startup DeepMind For More Than $500M

Posted Jan 26, 2014 by Catherine Shu (@catherineshu)

Zürcher Hochschule
für Angewandte Wissenschaften
zh
aw

Google will bu
reports that th
in talks to buy
couldn't disclose deal terms.

The acquisition was originally confirmed by Google to Re/code.

Google Acquires Artificial Intelligence Startup DeepMind For More Than $500M

Posted Jan 26, 2014 by Catherine Shu (@catherineshu)

Zürcher Hochschule für Angewandte Wissenschaften

zh aw

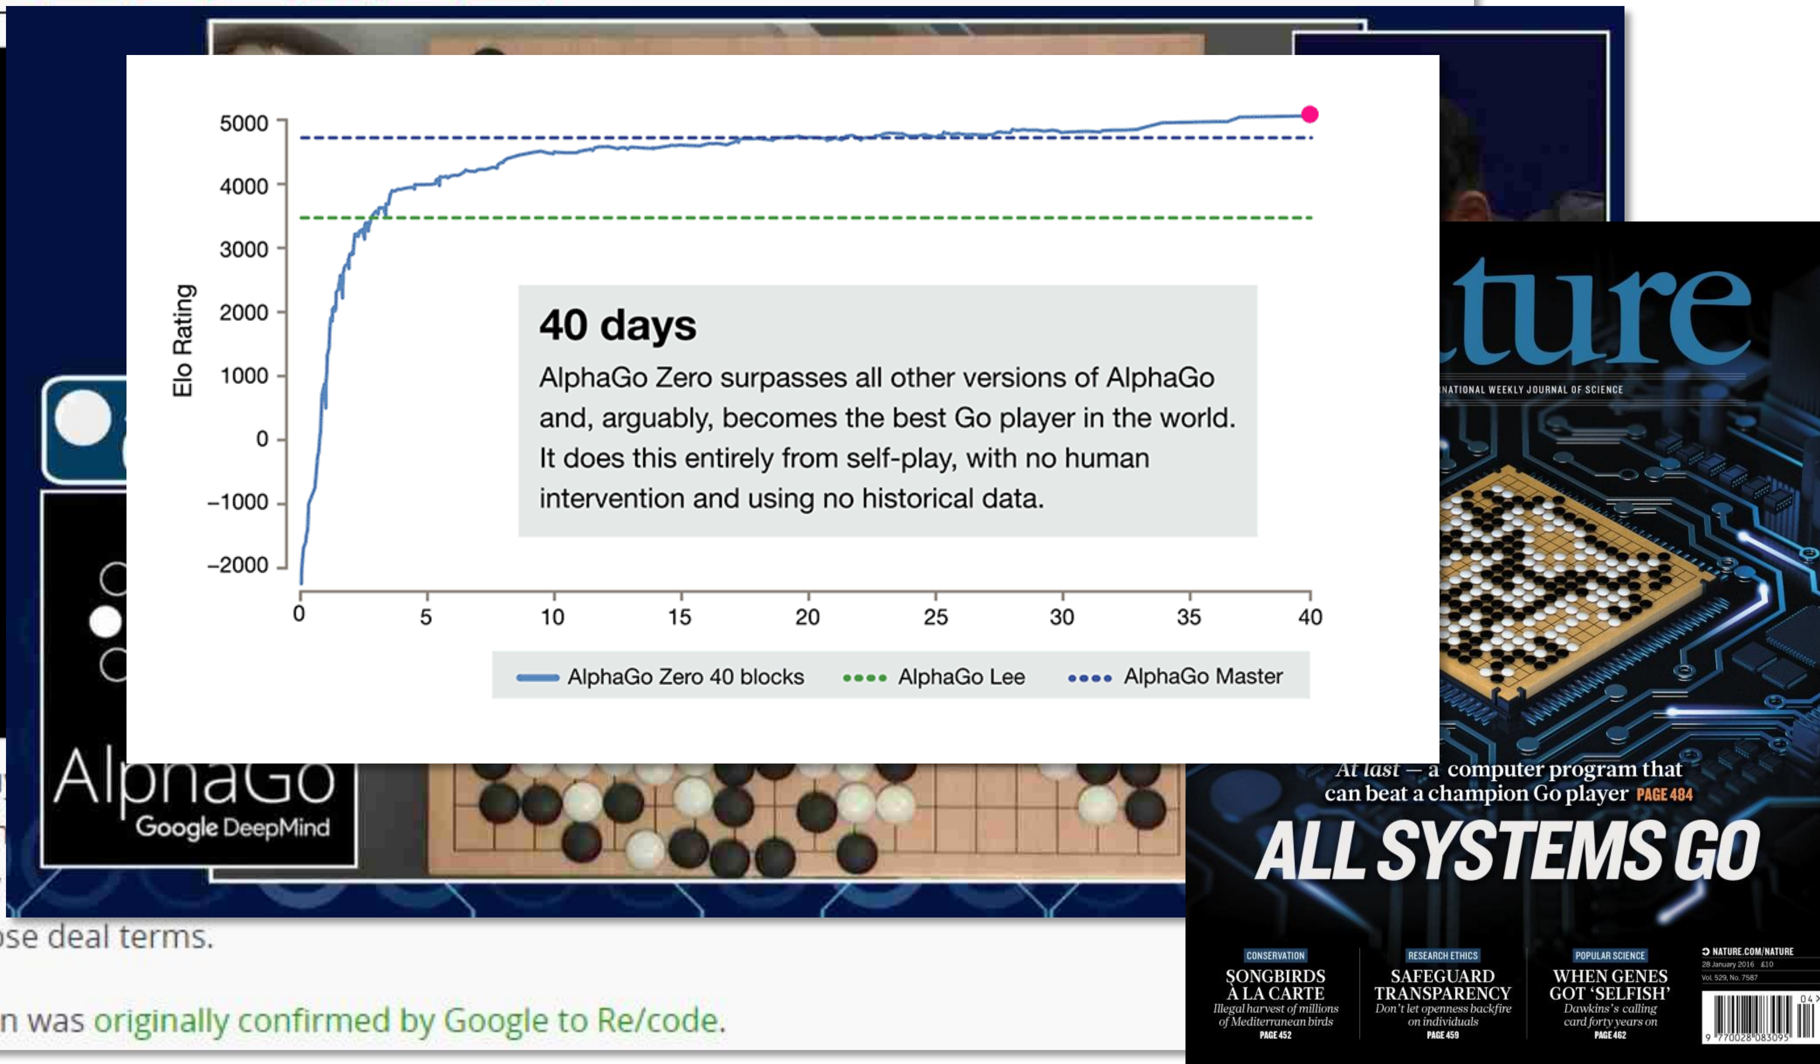

Google will bu

reports that th

in talks to buy

couldn't disclose deal terms.

The acquisition was originally confirmed by Google to Re/code.

Zürcher Hochschule
für Angewandte Wissenschaften

MIT Technology Review

Topics+ Top Stories

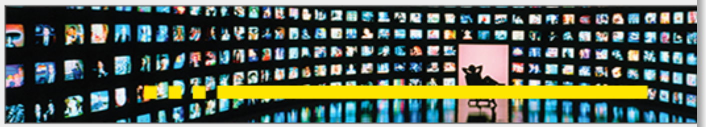

Computing

Algorithm Clones Van Gogh's Artistic Style and Pastes It onto Other Images, Movies

A deep neural network has learned to transfer artistic styles to other images.

by Emerging Technology from the arXiv May 10, 2016

The nature of artistic style is something of a mystery to most people. Think of Vincent Van Gogh's *Starry Night*, Picasso's work on cubism, or Edvard Munch's *The Scream*. All have a powerful, unique style that humans recognize easily.

MIT Technology Review

Topics+ Top Stories

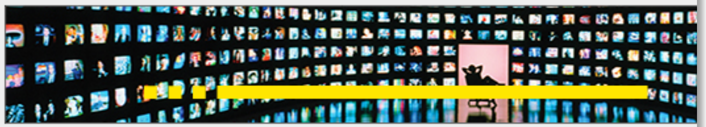

Computing

Algorithm Clones Van Gogh's Artistic Style and Pastes It onto Other Images, Movies

A deep neural network has learned to transfer artistic styles to other images.

by Emerging Technology from the arXiv May 10, 2016

f

The nature of artistic style is something of a mystery to most people. Think of Vincent Van Gogh's *Starr* ... n, or Edvard Munch's *The Scream* ... e th humans recognize easily.

Zürcher Hochschule
für Angewandte Wissenschaften

MIT Technology Review

Topics+ Top Stories

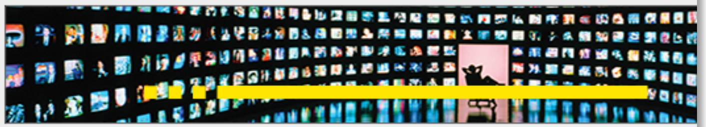

Computing

Algorithm Clones Van Gogh's Artistic Style and Pastes It onto Other Images, Movies

A deep neural network has learned to transfer artistic styles to other images.

by Emerging Technology from the arXiv May 10, 2016

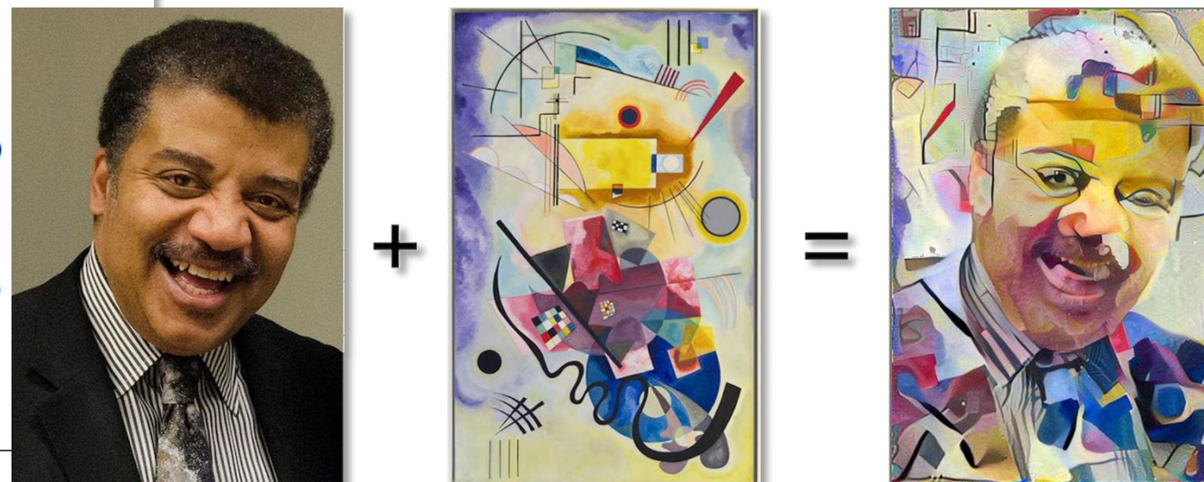

The nature of artistic style is something of a mystery to most people. Think of Vincent Van Gogh's *Starr*… n, or Edvard Munch's *The Scream*… e th… humans recognize easily.

MIT
Techn
Revie

Computing

Algorithm Clones Van Gogh's Artistic Style and Pastes It onto Other Images, Movies

A deep neural network has learned to transfer artistic styles to other images.

by Emerging Technology from the arXiv May 10, 2016

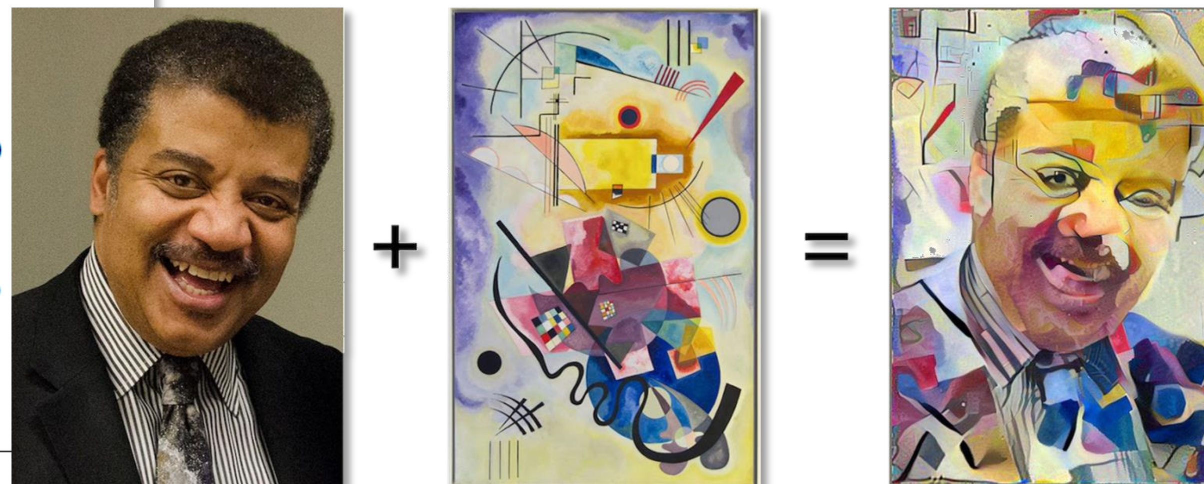

The nature of artistic style is something of a mystery to most people. Think
of Vincent Van Gogh's *Starr* ... n, or
Edvard Munch's *The Scream* ... e th
humans recognize easily.

MIT
Techn
Revie

Deep neural networks can now transfer the style of one photo onto another

And the results are impressive

by James Vincent | @jjvincent | Mar 30, 2017, 1:53pm EDT

Original photo Reference photo Result

You've probably heard of an AI technique known as "style transfer" — or, if you haven't heard of it, you've seen it. The process uses neural networks to apply the look and feel of one image to another, and appears in apps like Prisma and Facebook. These style transfers, however, are stylistic, not photorealistic. They look good because they look like they've been painted. Now a group of researchers from Cornell University and Adobe have augmented

NOW TRENDING

Computing

Algorith
Artistic
Other In

A deep neural n
other images.

by Emerging Tech

The nature of arti
of Vincent Van G
Edvard Munch's
humans recogni

Zürcher Hochschule
für Angewandte Wissenschaften

…und die Liste liesse sich fortsetzen!

…und die Liste liesse sich fortsetzen!


Brandon Amos About Blog

Image Completion with Deep Learning in TensorFlow

August 9, 2016

- Introduction
- Step 1: Interpreting images as samples from a probability distribution
  - How would you fill in the missing information?
  - But where does statistics fit in? These are images.
  - So how can we complete images?
- Step 2: Quickly generating fake images
  - Learning to generate new samples from an unknown probability distribution
  - [ML-Heavy] Generative Adversarial Net (GAN) building blocks
  - Using $G(z)$ to produce fake images
  - [ML-Heavy] Training DCGANs
  - Existing GAN
  - [ML-Heavy]
  - Running DCG
- Step 3: Finding the
  - Image compl
  - [ML-Heavy]
  - [ML-Heavy]
  - Completing y
- Conclusion
- Partial bibliography
- Bonus: Incomplete

Introduction

Content-aware fill is a p
completion and inpainti
do content-aware fill, im
"Semantic Image Inpaint
shows how to use deep l
some deeper portions fo
section can be skipped i
from images of faces. I h
completion.tensorflow.

We'll approach image co

1. We'll first interpret
2. This interpretation
3. Then we'll find the

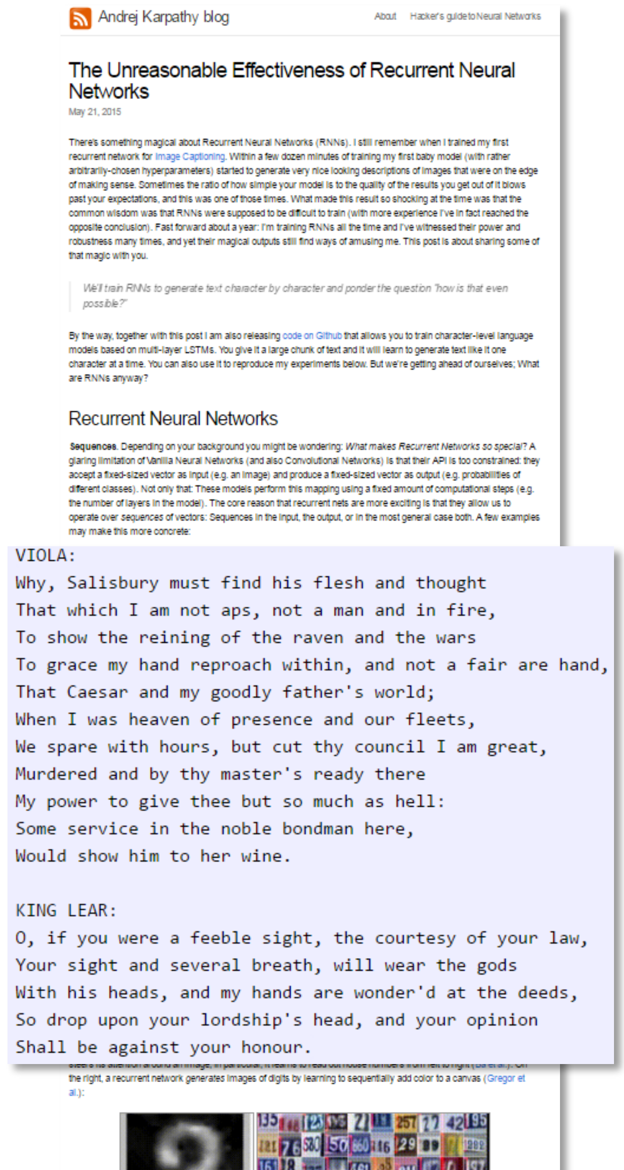
Andrej Karpathy blog About Hacker's guide to Neural Networks

The Unreasonable Effectiveness of Recurrent Neural Networks

May 21, 2015

There's something magical about Recurrent Neural Networks (RNNs). I still remember when I trained my first recurrent network for Image Captioning. Within a few dozen minutes of training my first baby model (with rather arbitrarily-chosen hyperparameters) started to generate very nice looking descriptions of images that were on the edge of making sense. Sometimes the ratio of how simple your model is to the quality of the results you get out of it blows past your expectations, and this was one of those times. What made this result so shocking at the time was that the common wisdom was that RNNs were supposed to be difficult to train (with more experience I've in fact reached the opposite conclusion). Fast forward about a year: I'm training RNNs all the time and I've witnessed their power and robustness many times, and yet their magical outputs still find ways of amusing me. This post is about sharing some of that magic with you.

> *We'll train RNNs to generate text character by character and ponder the question "how is that even possible?"*

By the way, together with this post I am also releasing code on Github that allows you to train character-level language models based on multi-layer LSTMs. You give it a large chunk of text and it will learn to generate text like it one character at a time. You can also use it to reproduce my experiments below. But we're getting ahead of ourselves; What are RNNs anyway?

Recurrent Neural Networks

Sequences. Depending on your background you might be wondering: *What makes Recurrent Networks so special?* A glaring limitation of Vanilla Neural Networks (and also Convolutional Networks) is that their API is too constrained: they accept a fixed-sized vector as input (e.g. an image) and produce a fixed-sized vector as output (e.g. probabilities of different classes). Not only that: These models perform this mapping using a fixed amount of computational steps (e.g. the number of layers in the model). The core reason that recurrent nets are more exciting is that they allow us to operate over *sequences* of vectors: Sequences in the input, the output, or in the most general case both. A few examples may make this more concrete:

```
VIOLA:
Why, Salisbury must find his flesh and thought
That which I am not aps, not a man and in fire,
To show the reining of the raven and the wars
To grace my hand reproach within, and not a fair are hand,
That Caesar and my goodly father's world;
When I was heaven of presence and our fleets,
We spare with hours, but cut thy council I am great,
Murdered and by thy master's ready there
My power to give thee but so much as hell:
Some service in the noble bondman here,
Would show him to her wine.

KING LEAR:
O, if you were a feeble sight, the courtesy of your law,
Your sight and several breath, will wear the gods
With his heads, and my hands are wonder'd at the deeds,
So drop upon your lordship's head, and your opinion
Shall be against your honour.
```

the right, a recurrent network generates images of digits by learning to sequentially add color to a canvas (Gregor et al.):

Zürcher Hochschule
für Angewandte Wissenschaften

…und die Liste liesse sich fortsetzen!

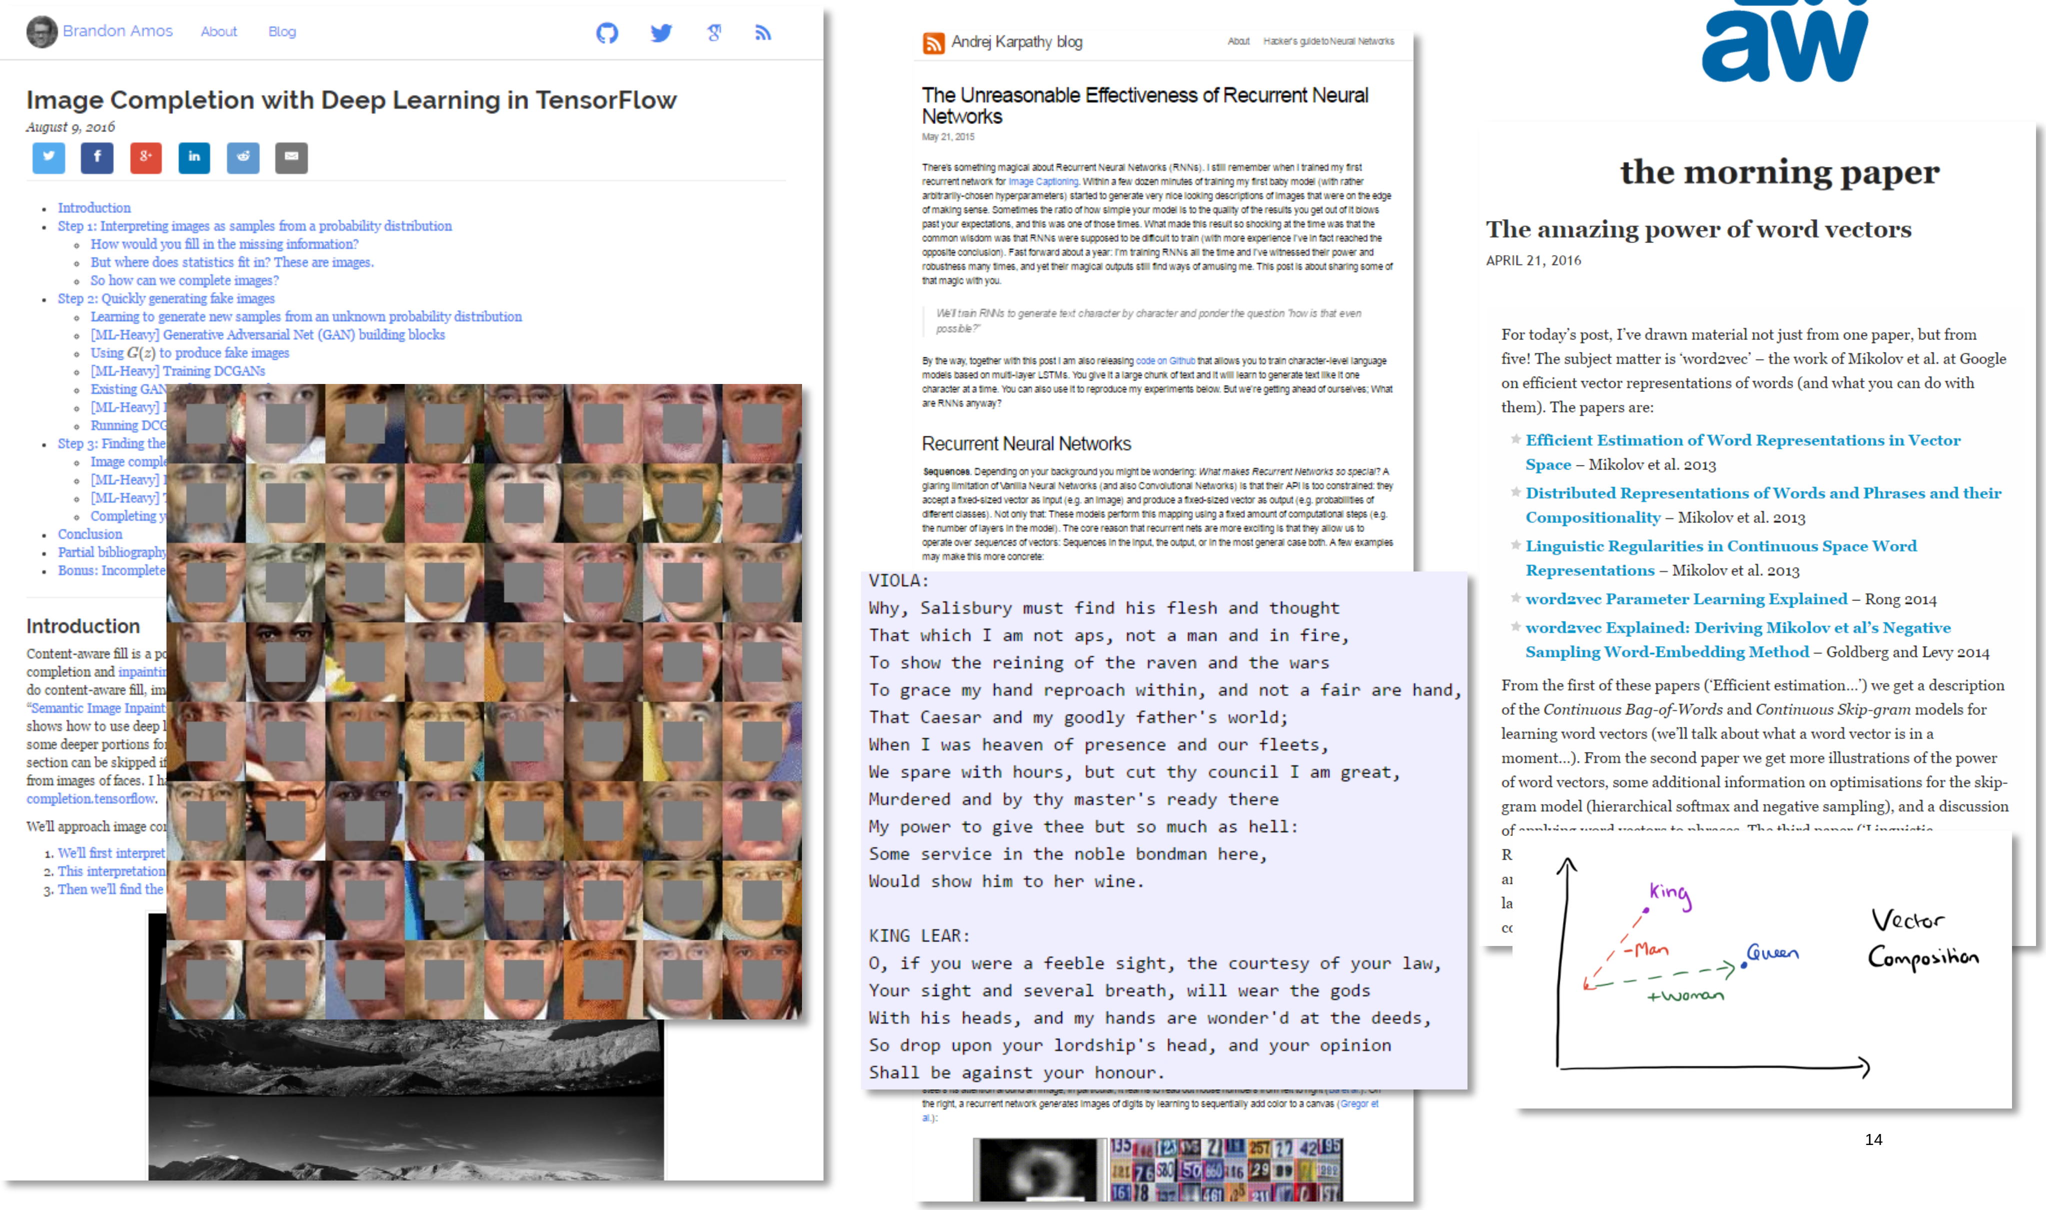

…und die Liste liesse sich fortsetzen!

Zürcher Hochschule
für Angewandte Wissenschaften

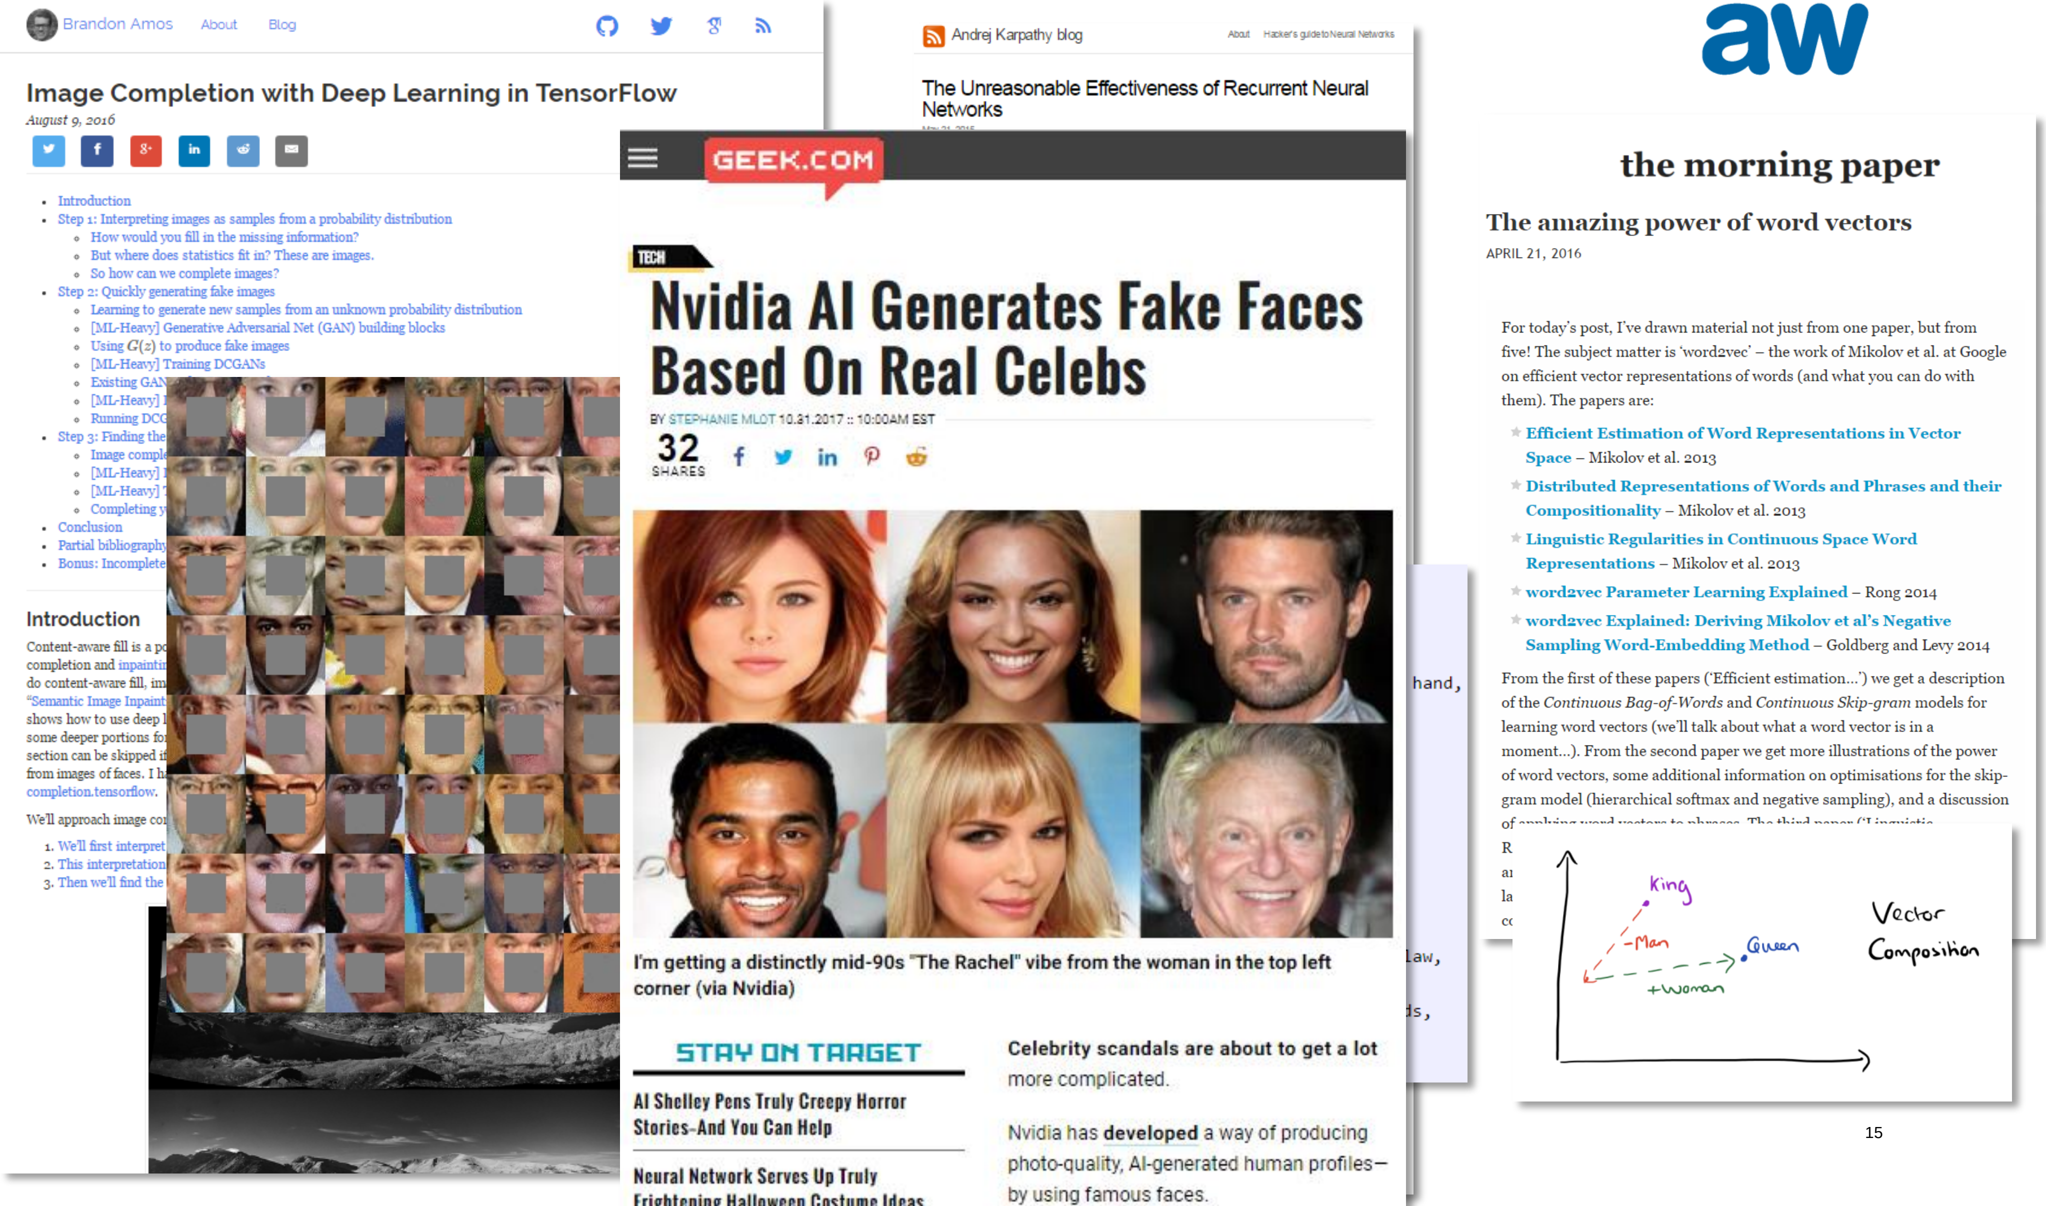

Was ist passiert?
Der ImageNet Wettbewerb

IMAGENET

1000 Kategorien
1 Mio. Beispiele

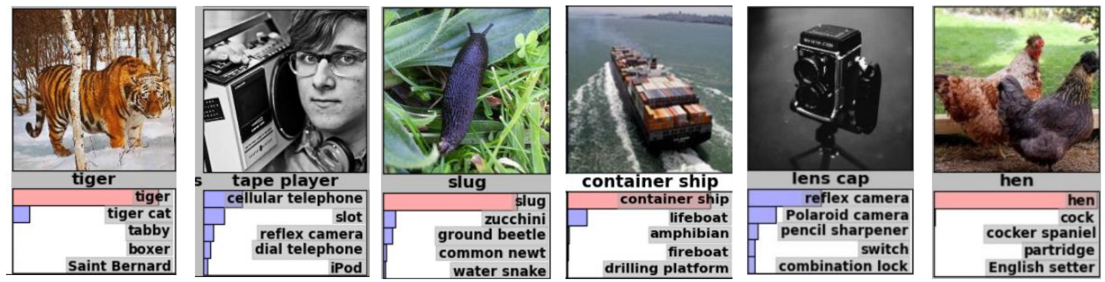

…

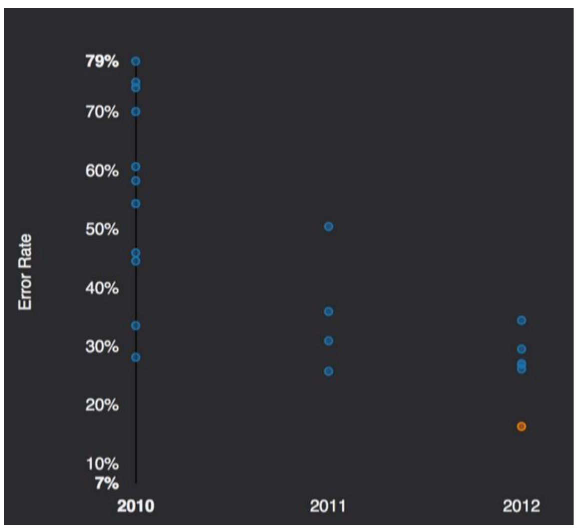

Was ist passiert?
Der ImageNet Wettbewerb

1000 Kategorien
1 Mio. Beispiele

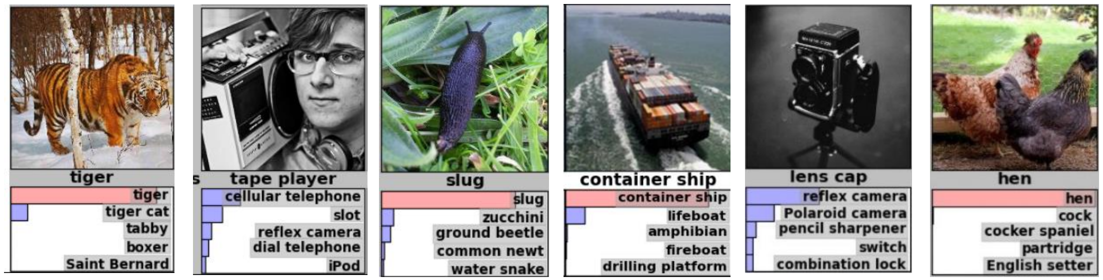

…

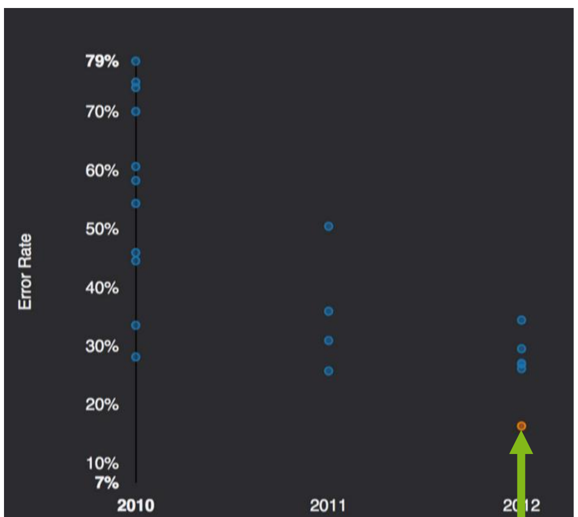

A. Krizhevsky verwendet als erster ein sog. «Deep Neural Network» (CNN)

Was ist passiert?
Der ImageNet Wettbewerb

1000 Kategorien
1 Mio. Beispiele

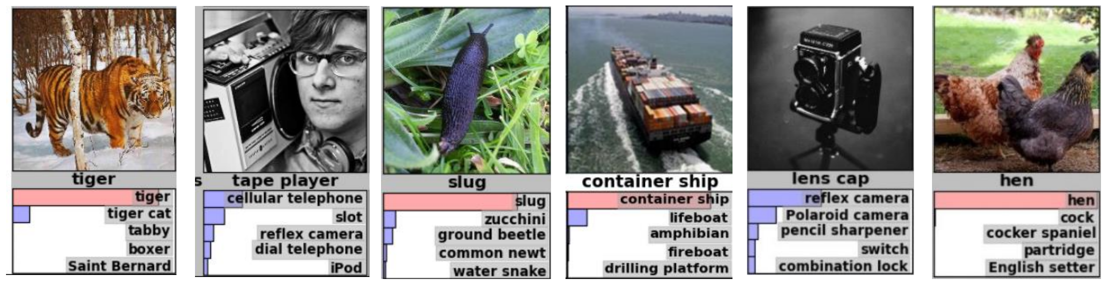

…

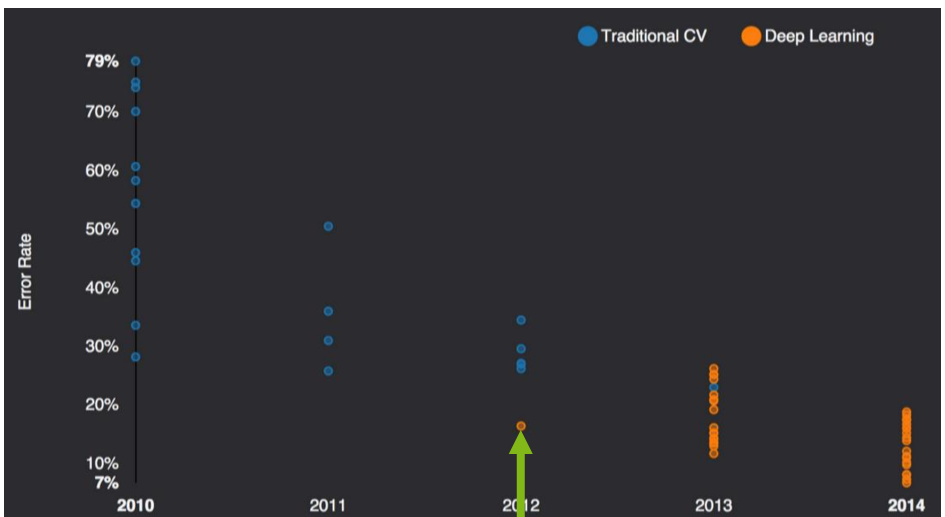

A. Krizhevsky verwendet als erster ein sog. «Deep Neural Network» (CNN)

Was ist passiert?
Der ImageNet Wettbewerb

1000 Kategorien
1 Mio. Beispiele

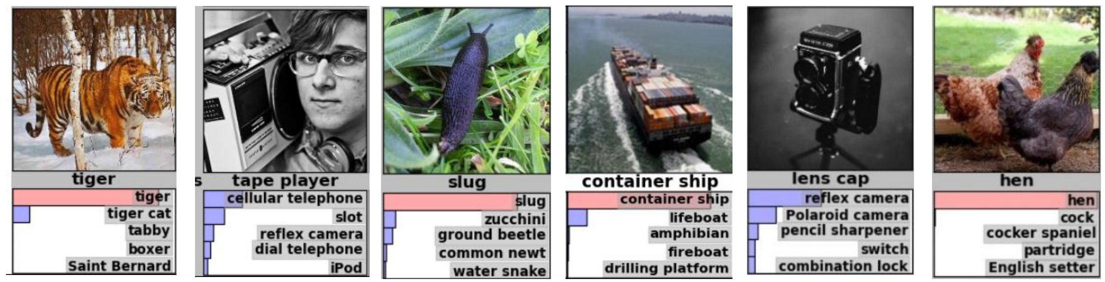

…

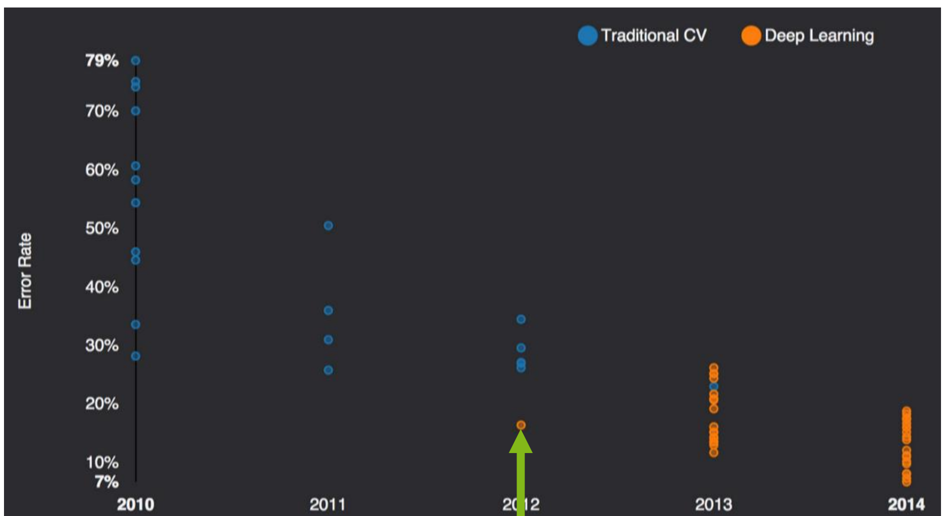

A. Krizhevsky verwendet als erster ein sog. «Deep Neural Network» (CNN)

2015: Computer *haben “Sehen” gelernt*

4.95% Microsoft (06. Februar)
→ Besser als Menschen (5.10%)

4.80% Google (11. Februar)

4.58% Baidu (11. Mai)

3.57% Microsoft (10. Dezember)

Zürcher Hochschule
für Angewandte Wissenschaften

2

Wie geht das?

Grundlage

Induktives überwachtes Lernen

Annahme

- Ein an *genügend viele* Beispiele angepasstes Modell…
- …wird auch auf unbekannte Daten **generalisieren**

Methode

- **Suchen der Parameter einer gegebenen Funktion**…
- …so dass für alle Beispiele Eingabe (Bild) auf Ausgabe («Auto») abgebildet wird

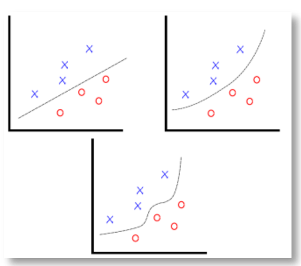

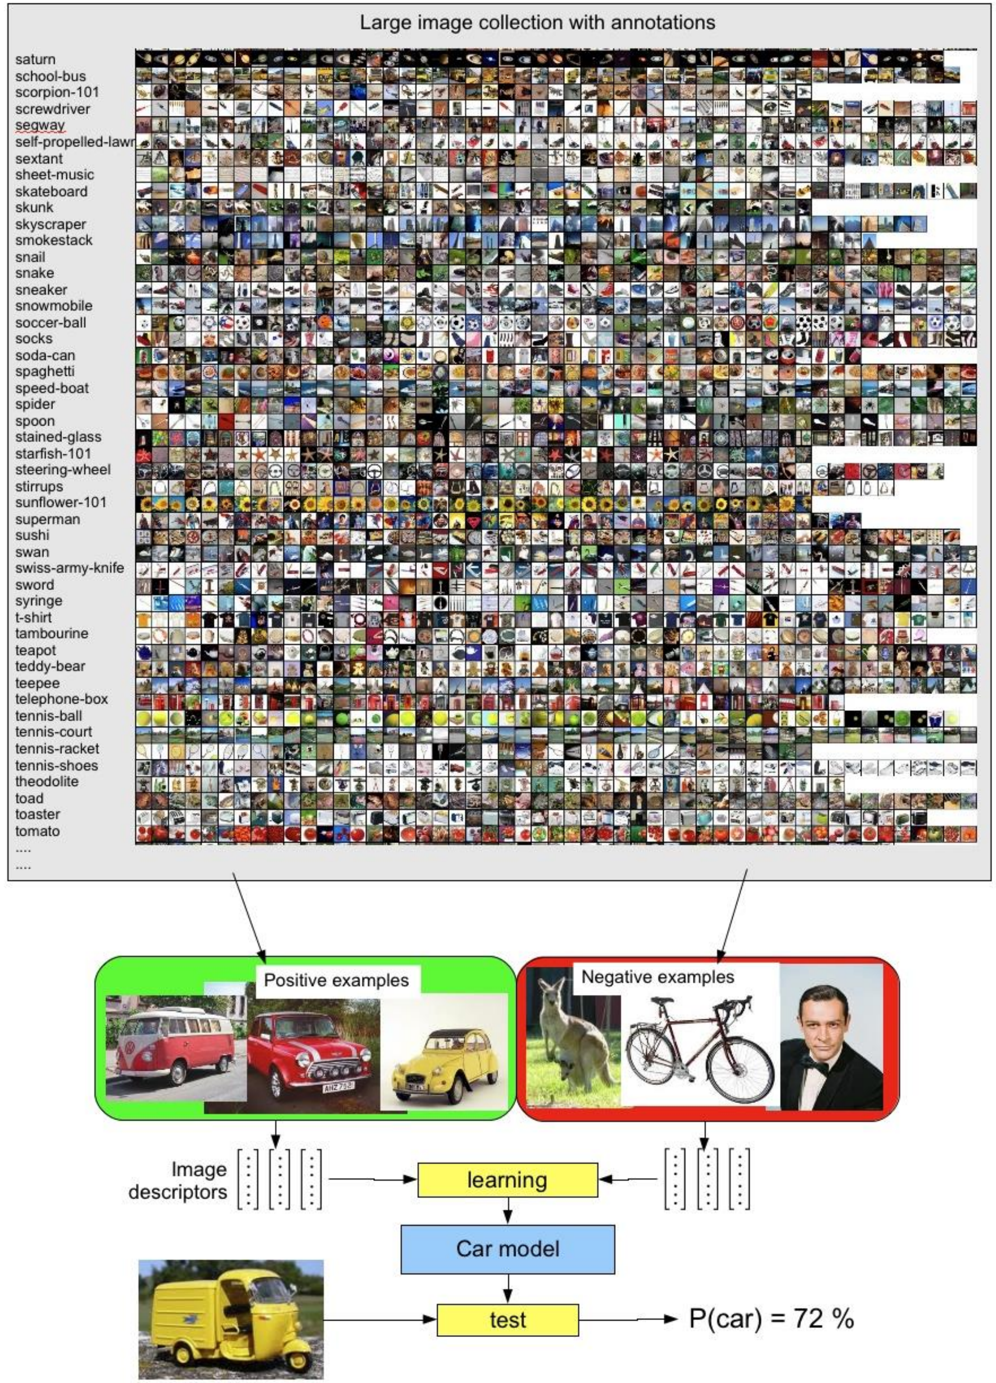

Quelle: http://lear.inrialpes.fr/job/postdoc-large-scale-classif-11-img/attribs_patchwork.jpg

Suche der Parameter *einer Funktion*?

Neuron

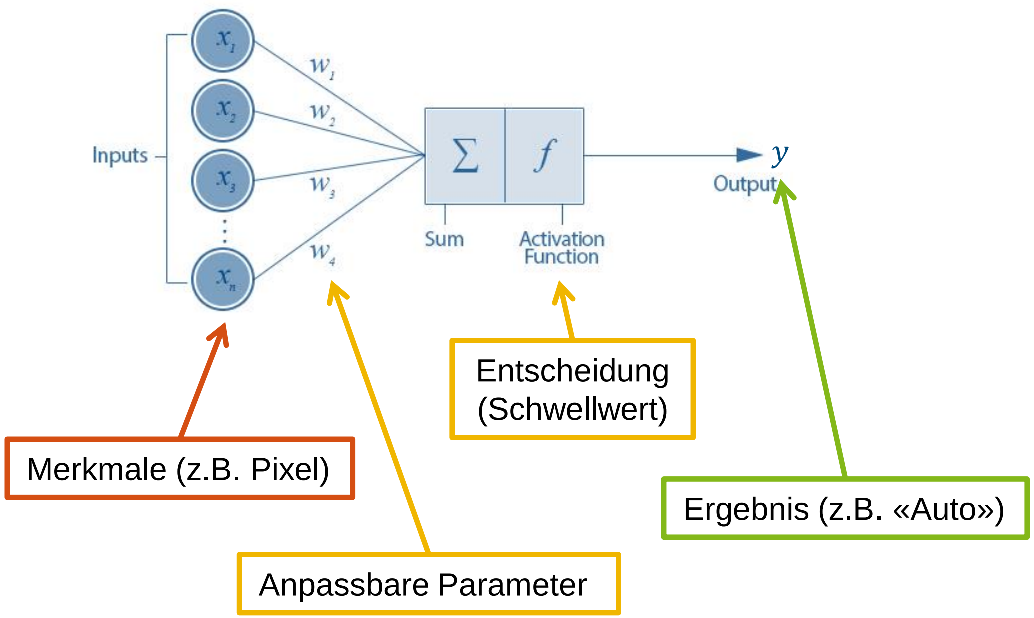

Suche der Parameter *einer Funktion*?

Neuron

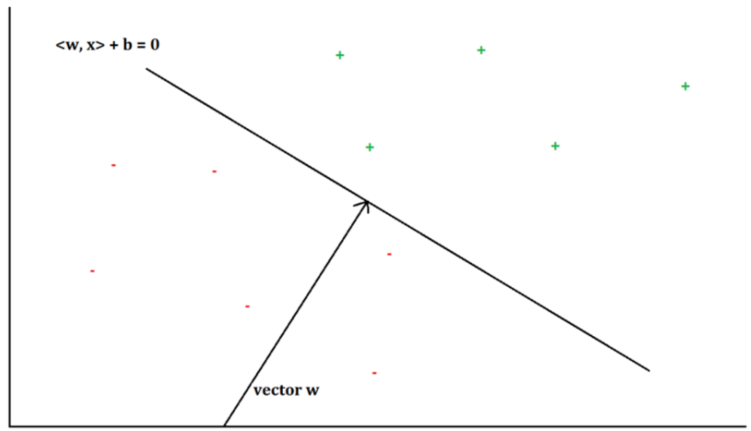

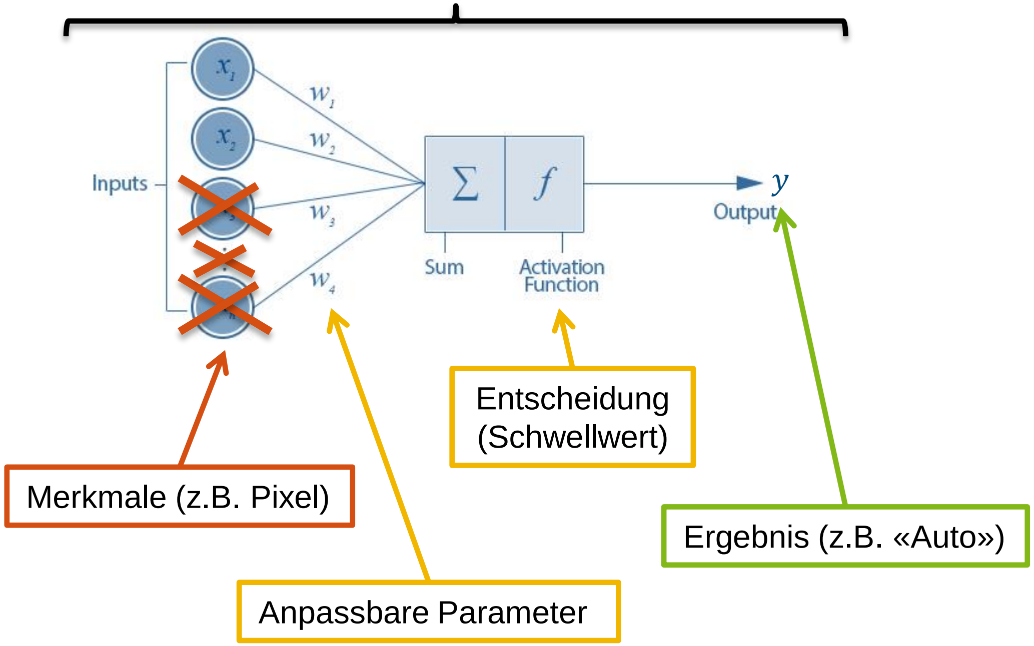

Zürcher Hochschule
für Angewandte Wissenschaften

Suche der Parameter *einer Funktion*?

Neuron

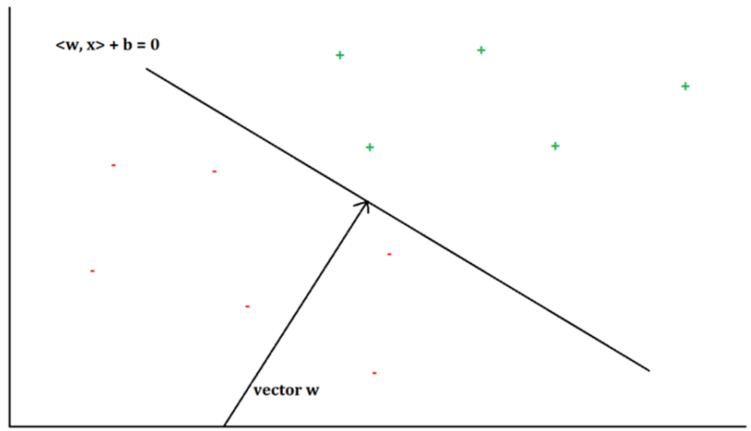

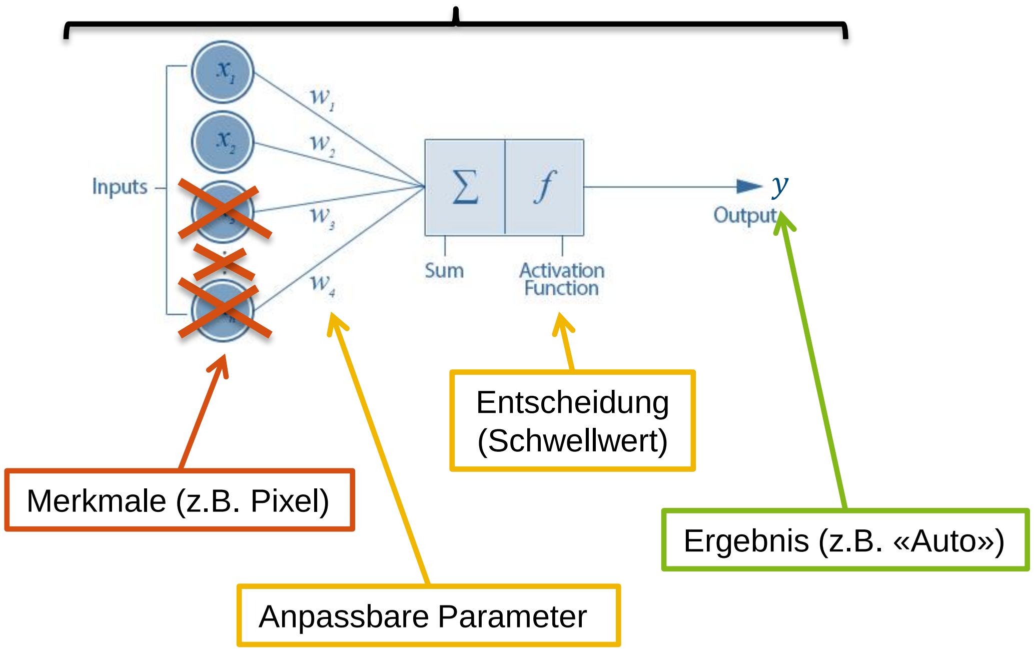

Neuronales Netz

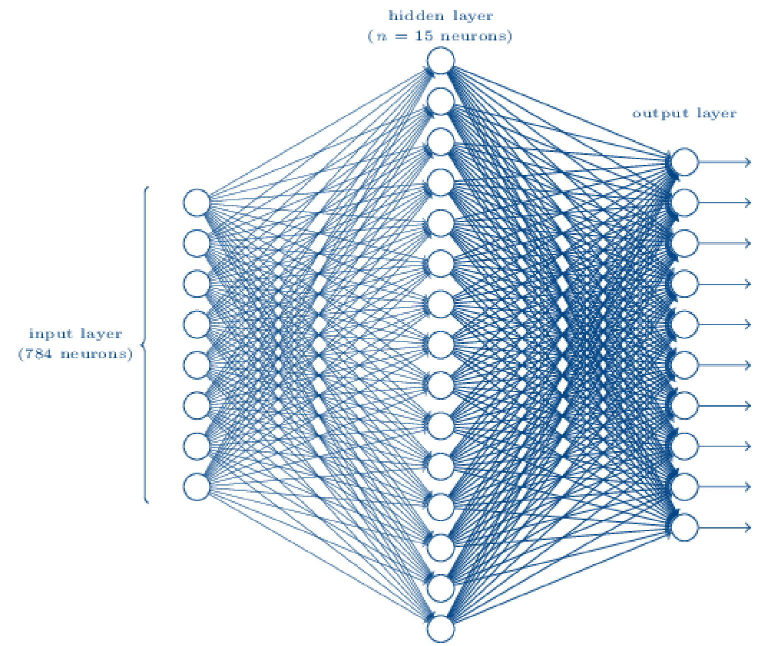

Suche der Parameter *einer Funktion*?

Neuron

Neuronales Netz

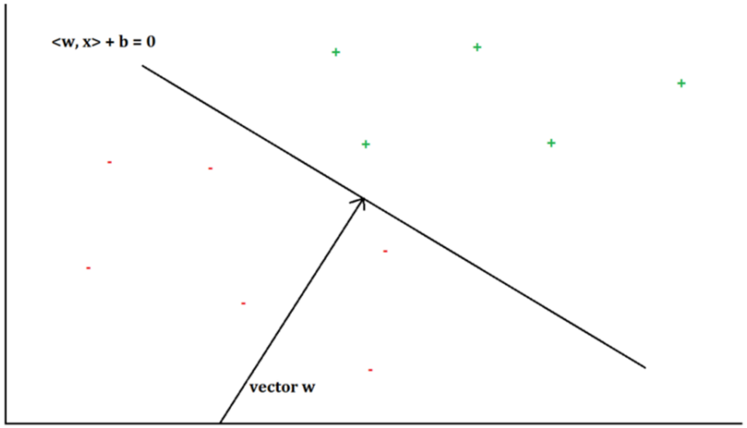

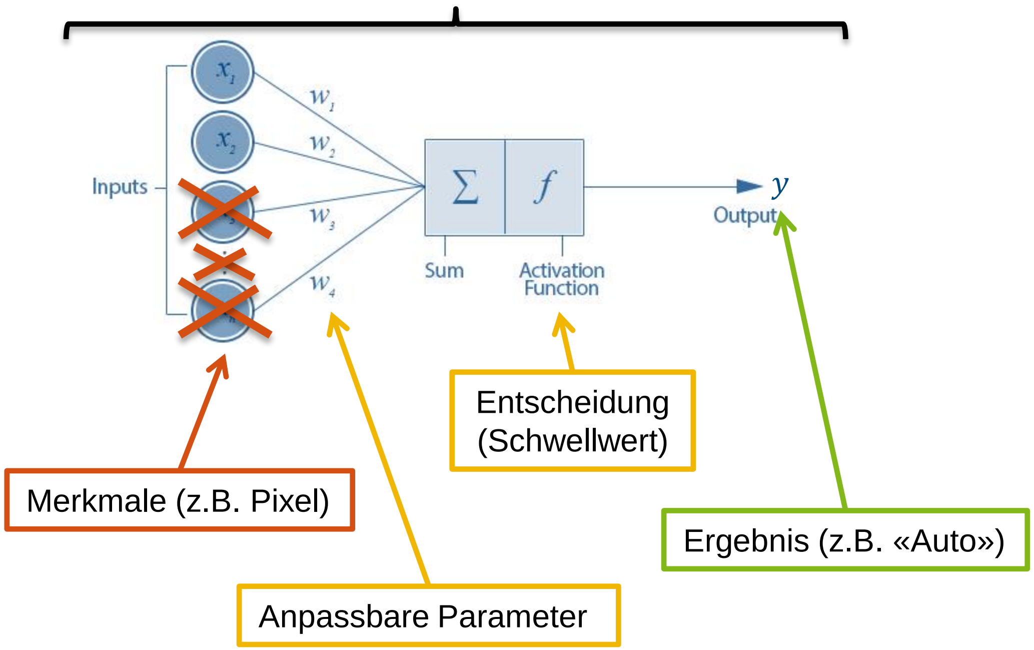

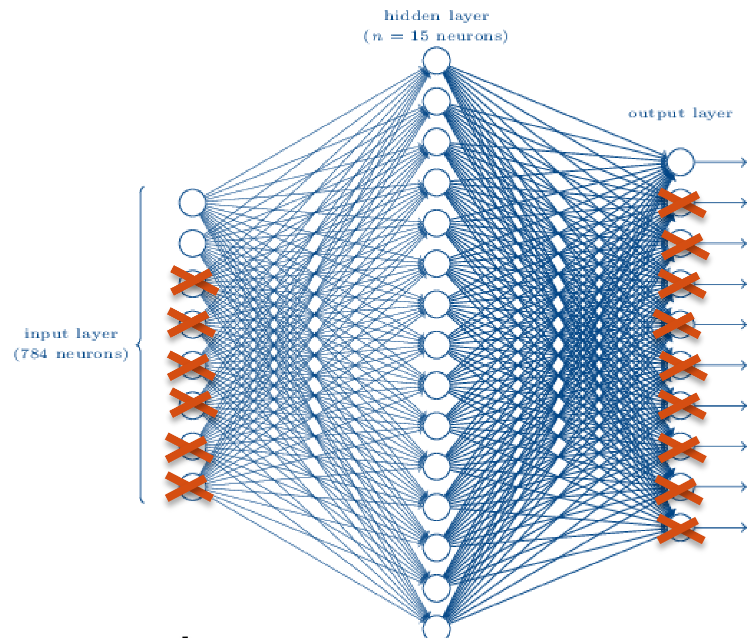

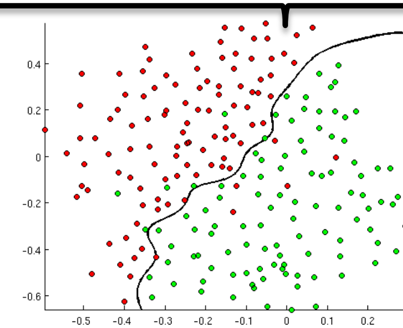

Suche der Parameter einer Funktion?

- Unser Neuronales Netz: $f_W(x) = y$
  mit Bild x, echtem Resultat y und Parametern $\boldsymbol{W}$
  ($\boldsymbol{W} = \{w_1, w_2\}$ anfangs zufällig gewählt)

- Fehlermass: $l(W) = \frac{1}{N}\sum_{i=1}^{N}(f_W(x_i) - y_i)^2$
  Durchschnitt der quadratischen Abweichungen über alle Bilder (Loss)

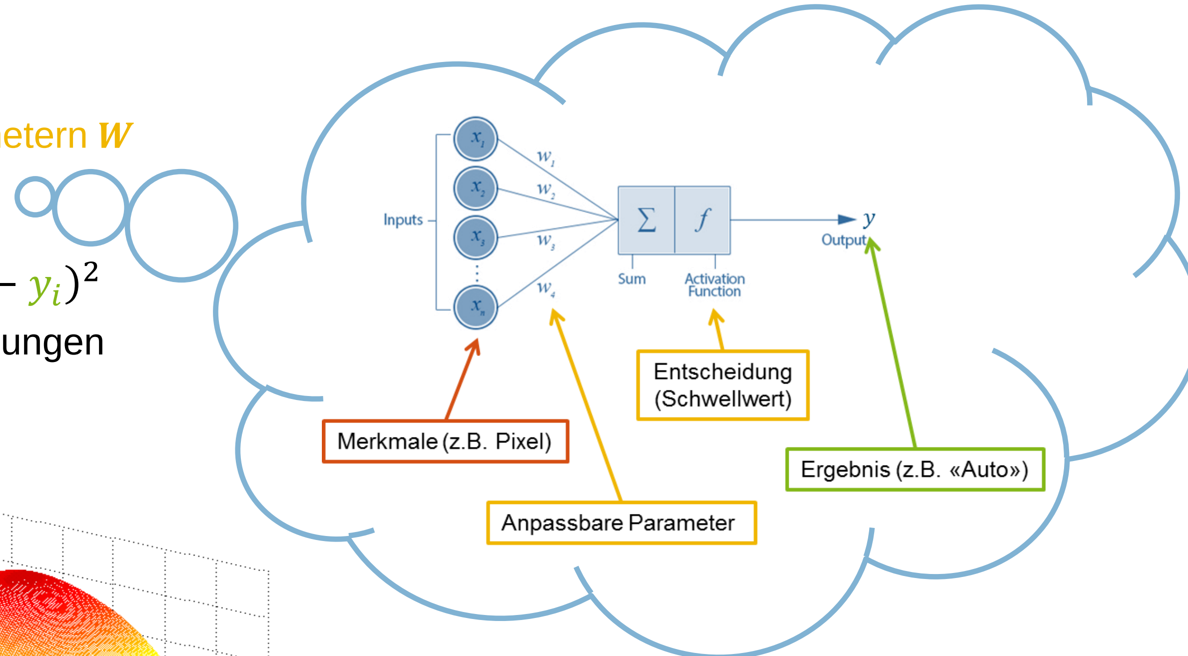

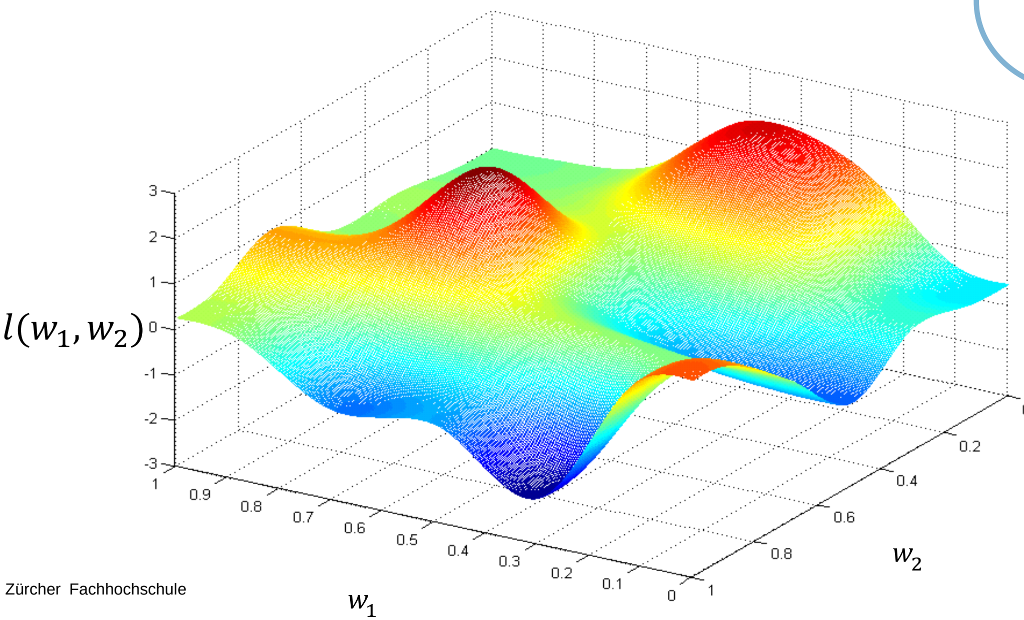

← Fehlerlandschaft

Suche der Parameter einer Funktion?

Wahrscheinlichkeit [%] für bestimmtes Ergebnis

- Unser Neuronales Netz: $f_W(x) = y$
  mit Bild x, echtem Resultat y und Parametern $\boldsymbol{W}$
  ($\boldsymbol{W} = \{w_1, w_2\}$ anfangs zufällig gewählt)

- Fehlermass: $l(W) = \frac{1}{N}\sum_{i=1}^{N}(f_W(x_i) - y_i)^2$
  Durchschnitt der quadratischen Abweichungen über alle Bilder (Loss)

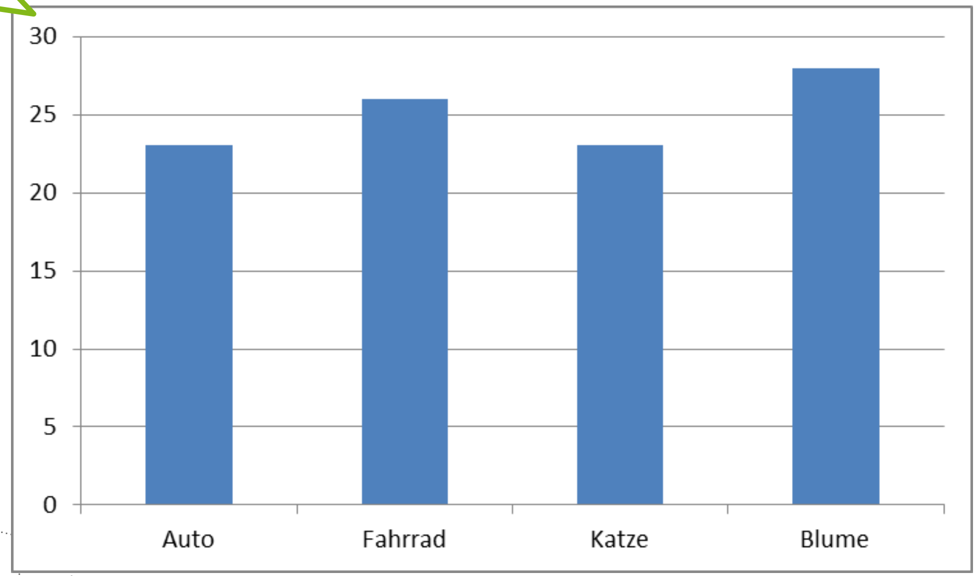

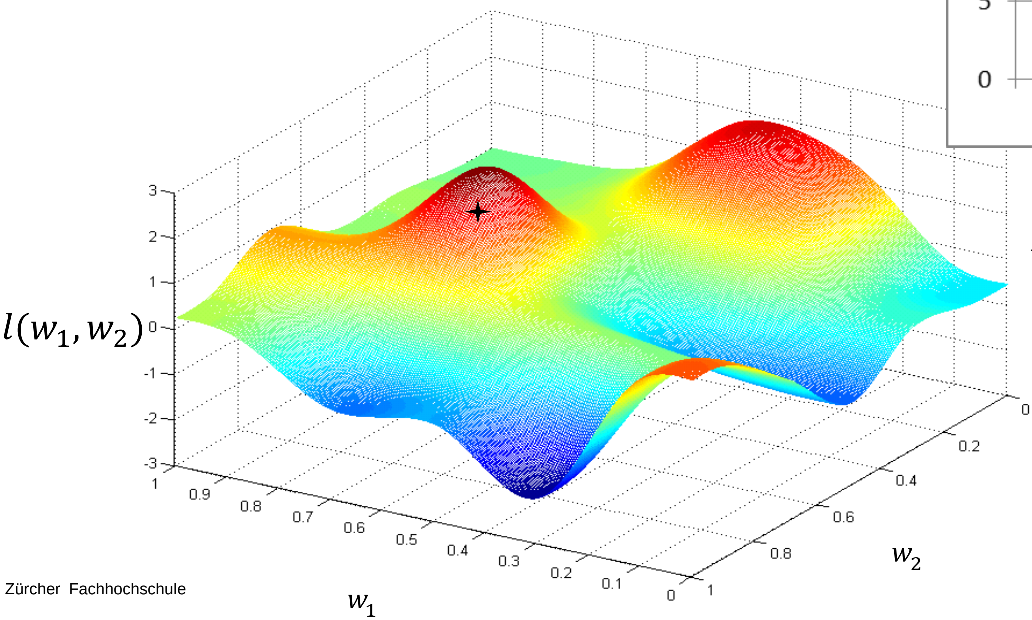

← Fehlerlandschaft

Suche der Parameter einer Funktion?

Wahrscheinlichkeit [%] für bestimmtes Ergebnis

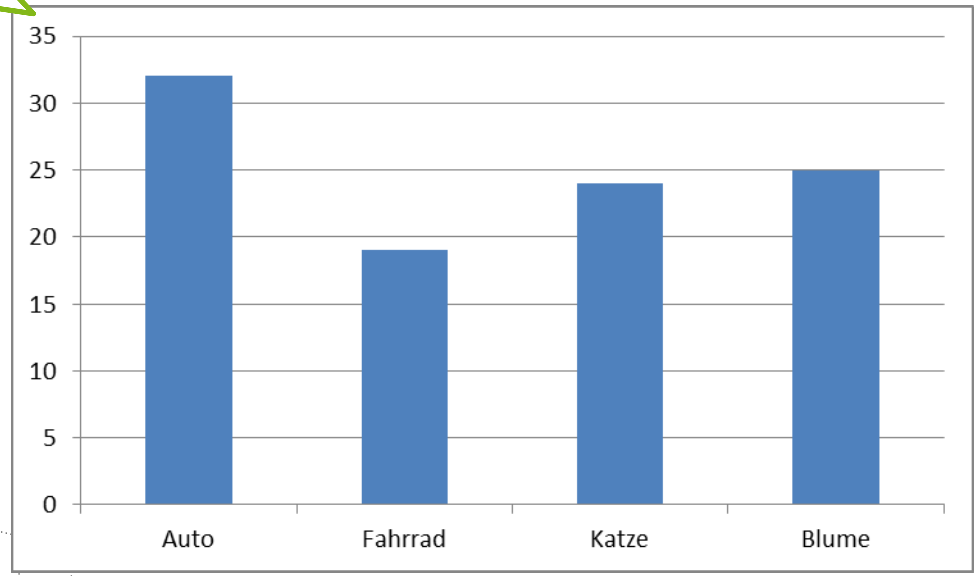

- Unser Neuronales Netz: $f_W(x) = y$
  mit Bild x, echtem Resultat y und Parametern $\boldsymbol{W}$
  ($\boldsymbol{W} = \{w_1, w_2\}$ anfangs zufällig gewählt)
- Fehlermass: $l(W) = \frac{1}{N}\sum_{i=1}^{N}(f_W(x_i) - y_i)^2$
  Durchschnitt der quadratischen Abweichungen über alle Bilder (Loss)

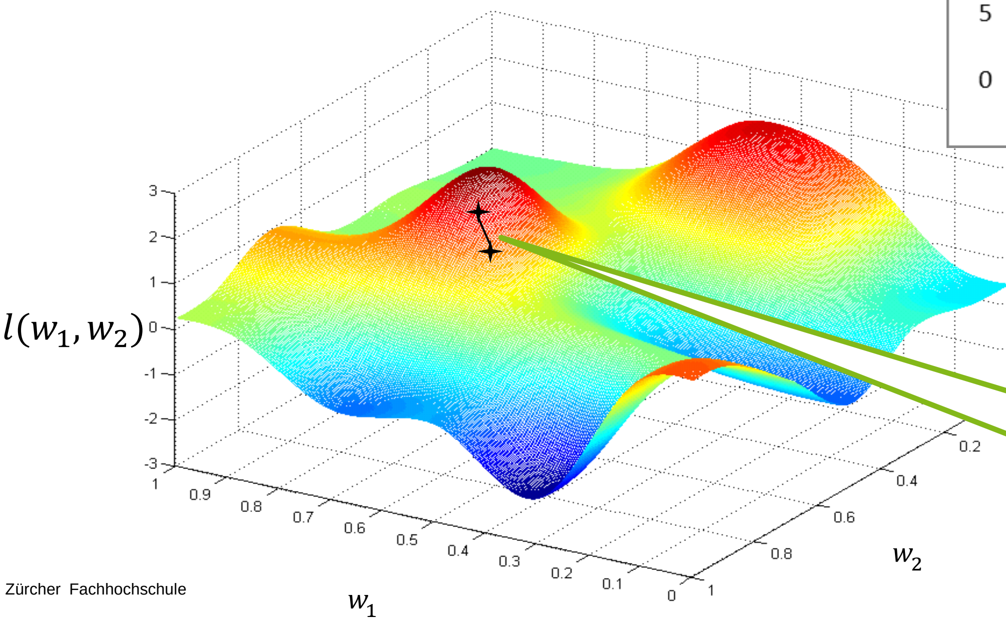

← Fehlerlandschaft

Methode: Anpassung der Gewichte von f in Richtung der steilsten Steigung (abwärts) von J

Suche der Parameter einer Funktion?

Wahrscheinlichkeit [%] für bestimmtes Ergebnis

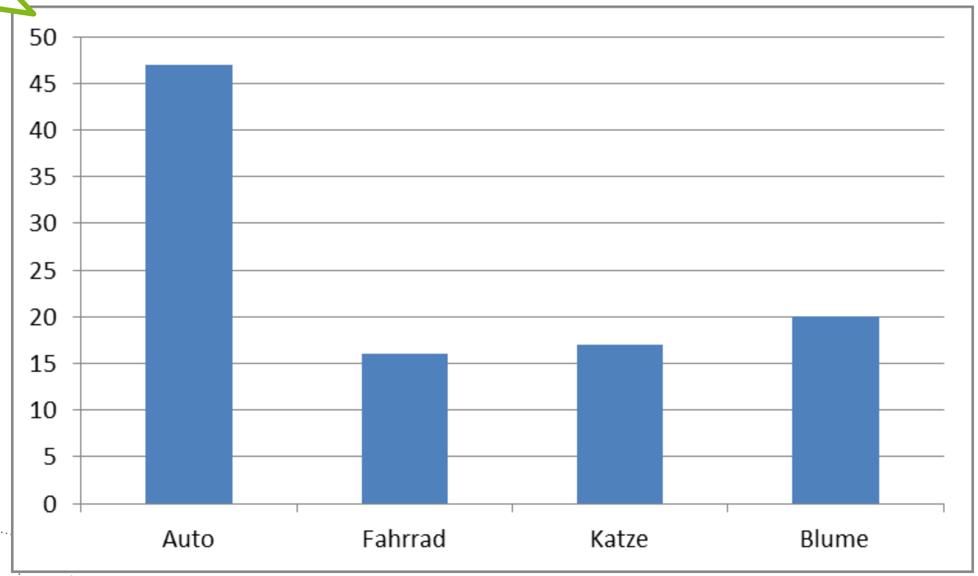

- Unser Neuronales Netz: $f_W(x) = y$
  mit Bild x, echtem Resultat y und Parametern $\boldsymbol{W}$
  ($\boldsymbol{W} = \{w_1, w_2\}$ anfangs zufällig gewählt)

- Fehlermass: $l(W) = \frac{1}{N}\sum_{i=1}^{N}(f_W(x_i) - y_i)^2$
  Durchschnitt der quadratischen Abweichungen über alle Bilder (Loss)

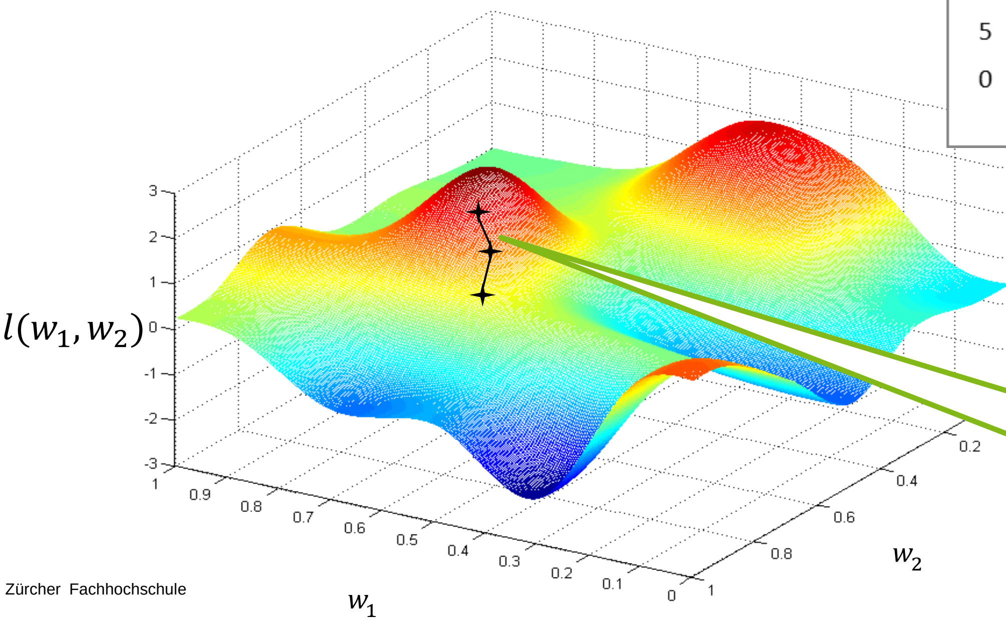

← Fehlerlandschaft

Methode: Anpassung der Gewichte von f in Richtung der steilsten Steigung (abwärts) von J

Suche der Parameter einer Funktion?

Wahrscheinlichkeit [%] für bestimmtes Ergebnis

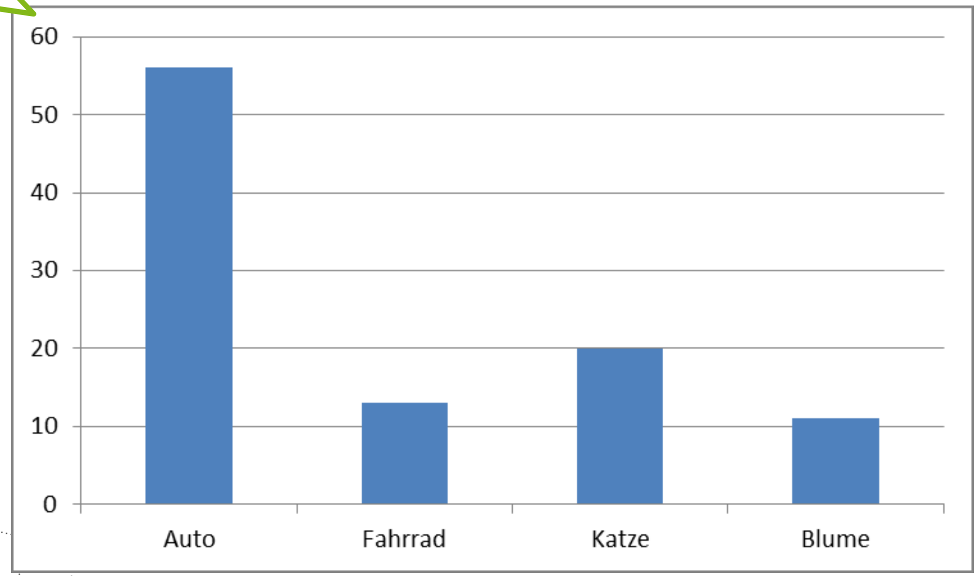

- Unser Neuronales Netz: $f_W(x) = y$
  mit Bild x, echtem Resultat y und Parametern $\boldsymbol{W}$
  ($\boldsymbol{W} = \{w_1, w_2\}$ anfangs zufällig gewählt)

- Fehlermass: $l(W) = \frac{1}{N}\sum_{i=1}^{N}(f_W(x_i) - y_i)^2$
  Durchschnitt der quadratischen Abweichungen über alle Bilder (Loss)

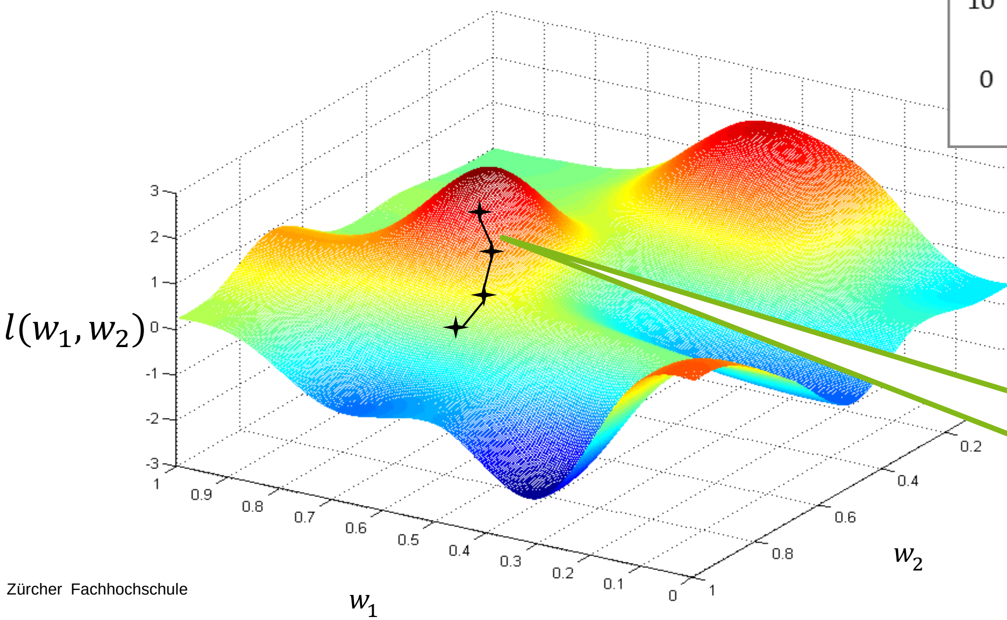

← Fehlerlandschaft

Methode: Anpassung der Gewichte von f in Richtung der steilsten Steigung (abwärts) von J

Suche der Parameter einer Funktion?

Wahrscheinlichkeit [%] für bestimmtes Ergebnis

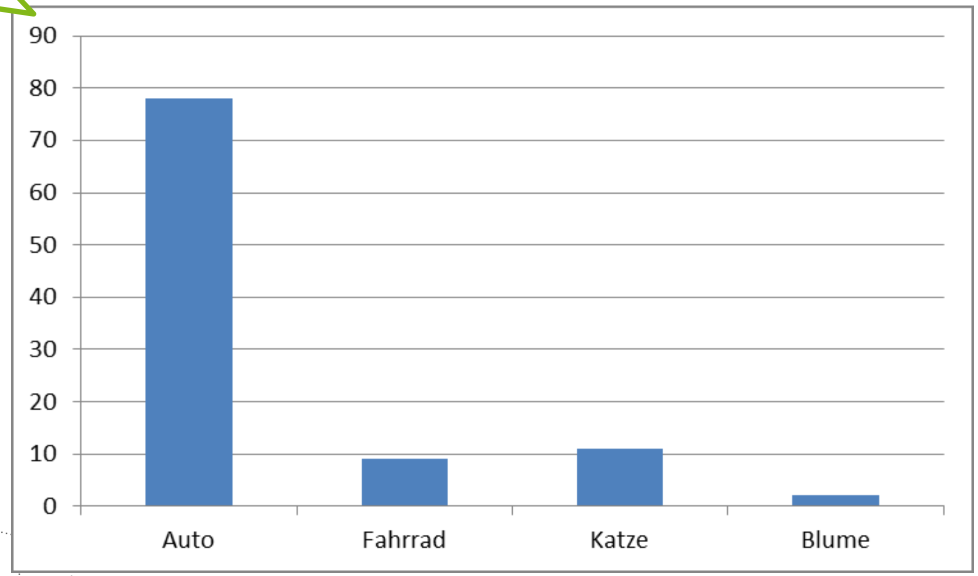

- Unser Neuronales Netz: $f_W(x) = y$
  mit Bild x, echtem Resultat y und Parametern $\boldsymbol{W}$
  ($\boldsymbol{W} = \{w_1, w_2\}$ anfangs zufällig gewählt)

- Fehlermass: $l(W) = \frac{1}{N}\sum_{i=1}^{N}(f_W(x_i) - y_i)^2$
  Durchschnitt der quadratischen Abweichungen über alle Bilder (Loss)

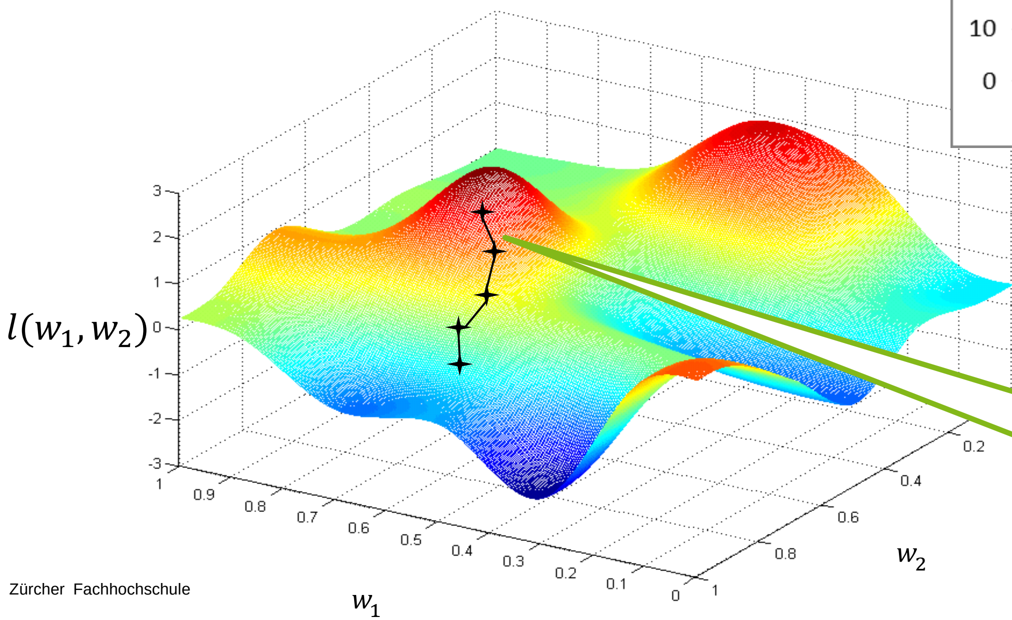

← Fehlerlandschaft

Methode: Anpassung der Gewichte von f in Richtung der steilsten Steigung (abwärts) von J

Suche der Parameter einer Funktion?

Wahrscheinlichkeit [%] für bestimmtes Ergebnis

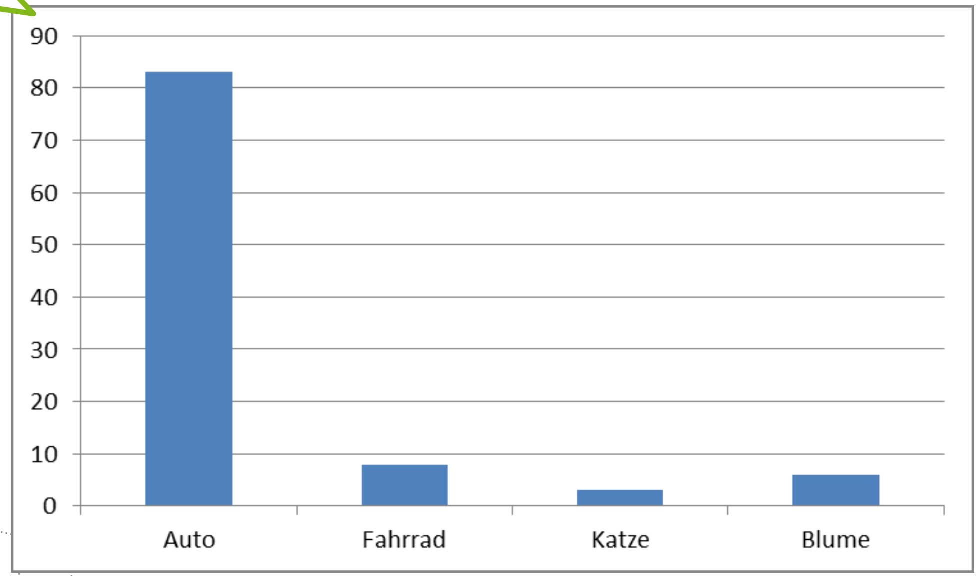

- Unser Neuronales Netz: $f_{W}(x) = y$
  mit Bild x, echtem Resultat y und Parametern $\boldsymbol{W}$
  ($\boldsymbol{W} = \{w_1, w_2\}$ anfangs zufällig gewählt)

- Fehlermass: $l(W) = \frac{1}{N}\sum_{i=1}^{N}(f_{W}(x_i) - y_i)^2$
  Durchschnitt der quadratischen Abweichungen über alle Bilder (Loss)

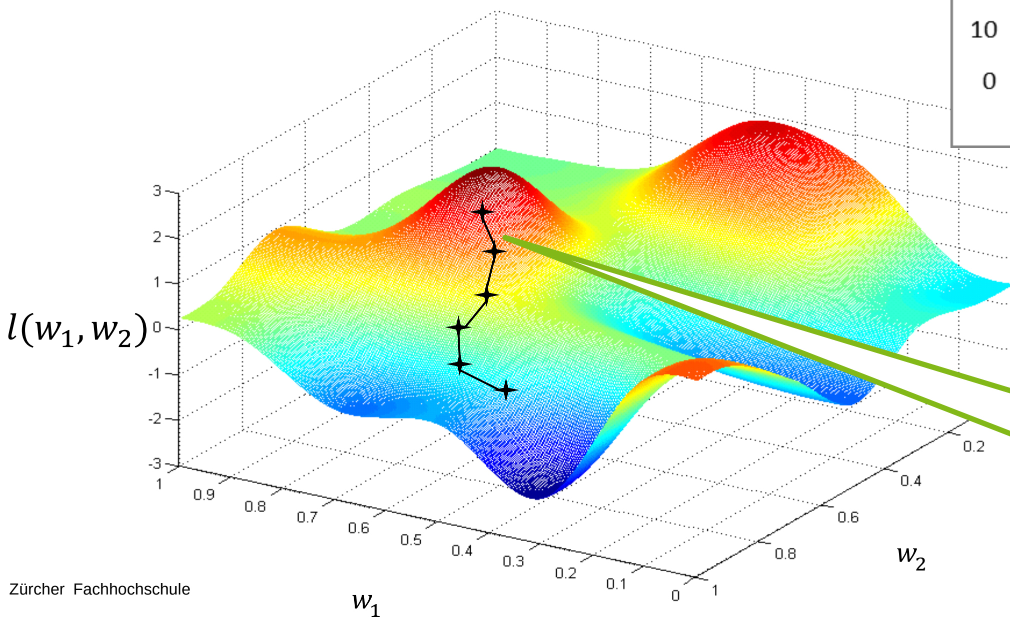

← Fehlerlandschaft

Methode: Anpassung der Gewichte von f in Richtung der steilsten Steigung (abwärts) von J

Suche der Parameter einer Funktion?

Wahrscheinlichkeit [%] für bestimmtes Ergebnis

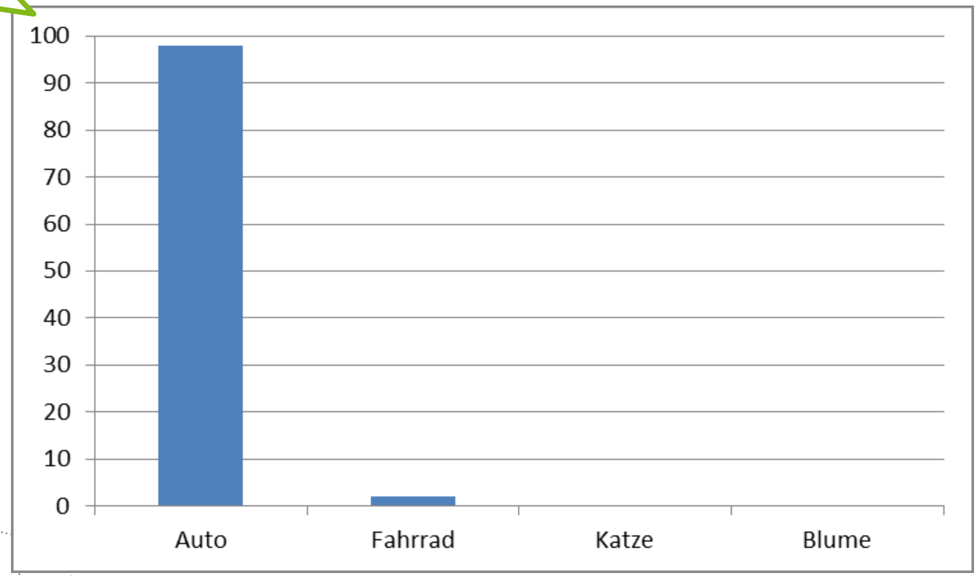

- Unser Neuronales Netz: $f_W(x) = y$
  mit Bild x, echtem Resultat y und Parametern $\boldsymbol{W}$
  ($\boldsymbol{W} = \{w_1, w_2\}$ anfangs zufällig gewählt)

- Fehlermass: $l(W) = \frac{1}{N}\sum_{i=1}^{N}(f_W(x_i) - y_i)^2$
  Durchschnitt der quadratischen Abweichungen über alle Bilder (Loss)

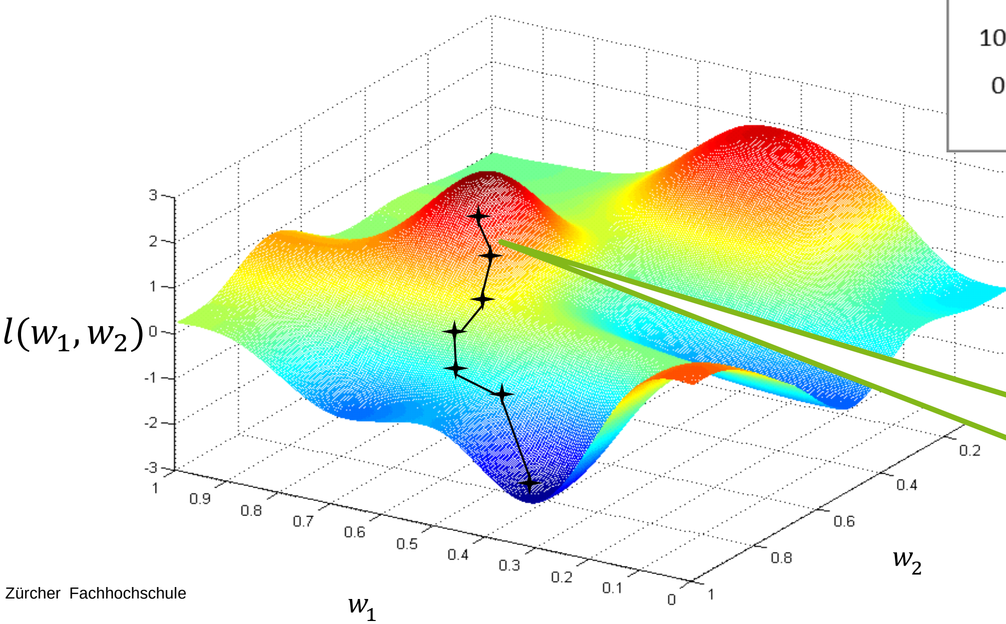

← Fehlerlandschaft

Methode: Anpassung der Gewichte von f in Richtung der steilsten Steigung (abwärts) von J

Was «sieht» das Neuronale Netz?
Hierarchien komplexer werdender Merkmale

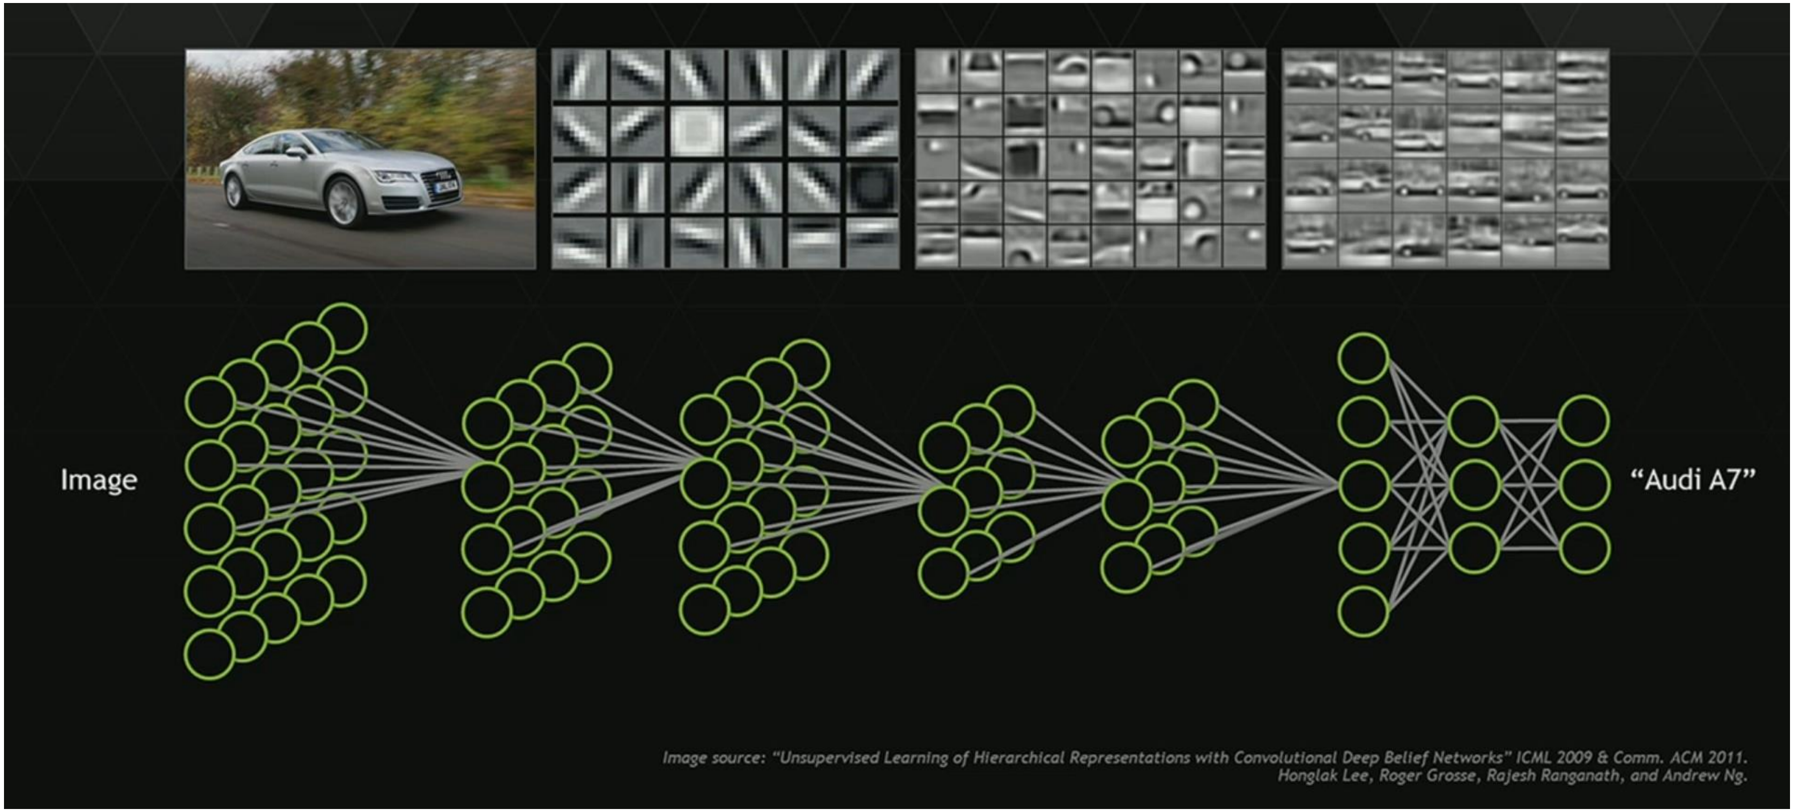

Quellen: https://www.pinterest.com/explore/artificial-neural-network/
Olah, et al., "Feature Visualization", Distill, 2017, https://distill.pub/2017/feature-visualization/.

Was «sieht» das Neuronale Netz?
Hierarchien komplexer werdender Merkmale

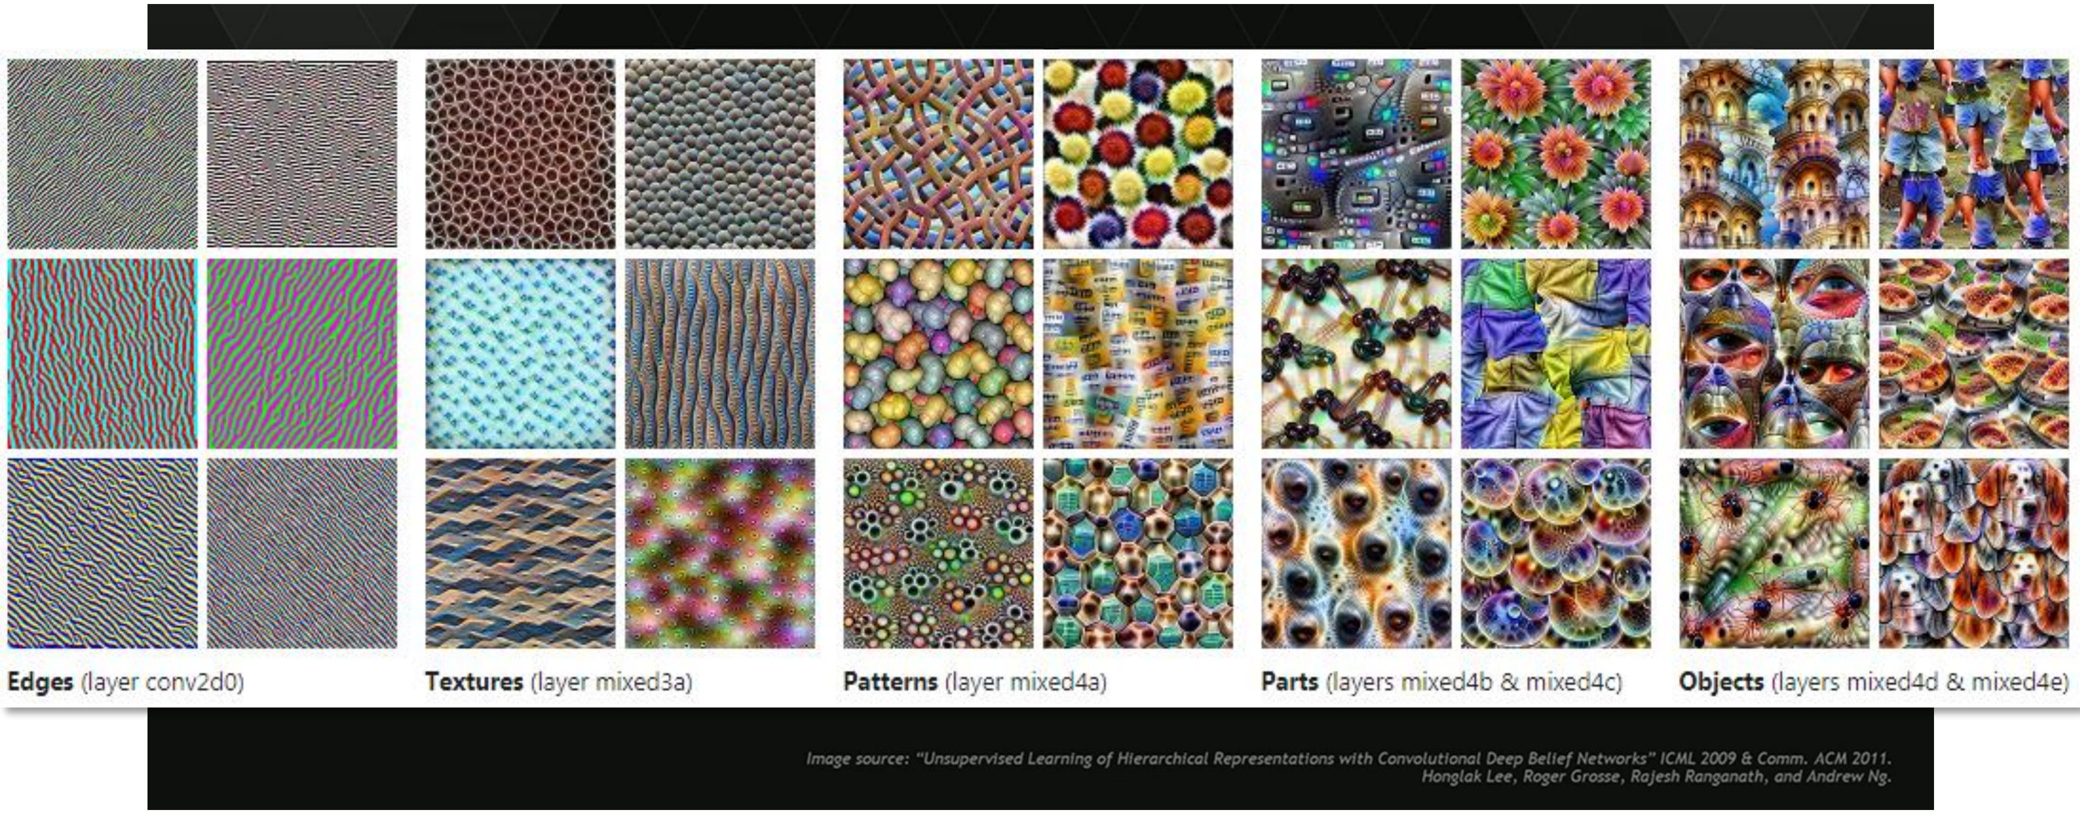

Quellen: https://www.pinterest.com/explore/artificial-neural-network/
Olah, et al., "Feature Visualization", Distill, 2017, https://distill.pub/2017/feature-visualization/.

Zürcher Hochschule
für Angewandte Wissenschaften

Wie schlussfolgert die Maschine?
«Debugging» für Einblicke in die vermeintliche «Black Box»

Verdeutlichen ein Problem:
- Adversarial Examples

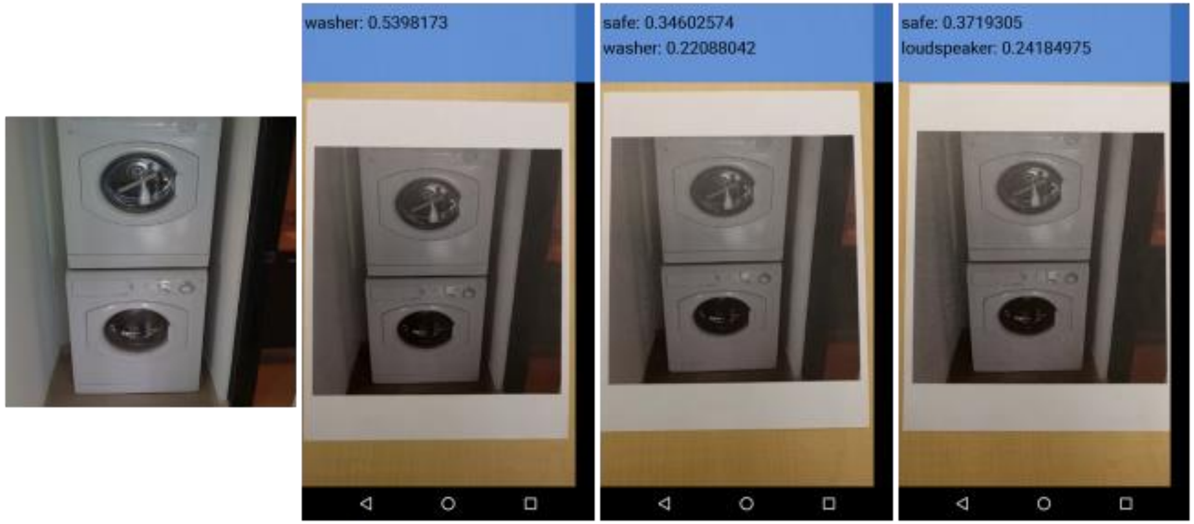

(a) Image from dataset (b) Clean image (c) Adv. image, $\epsilon = 4$ (d) Adv. image, $\epsilon = 8$

https://blog.openai.com/adversarial-example-research/

Wie schlussfolgert die Maschine?

«Debugging» für Einblicke in die vermeintliche «Black Box»

Verdeutlichen ein Problem:

- Adversarial Examples

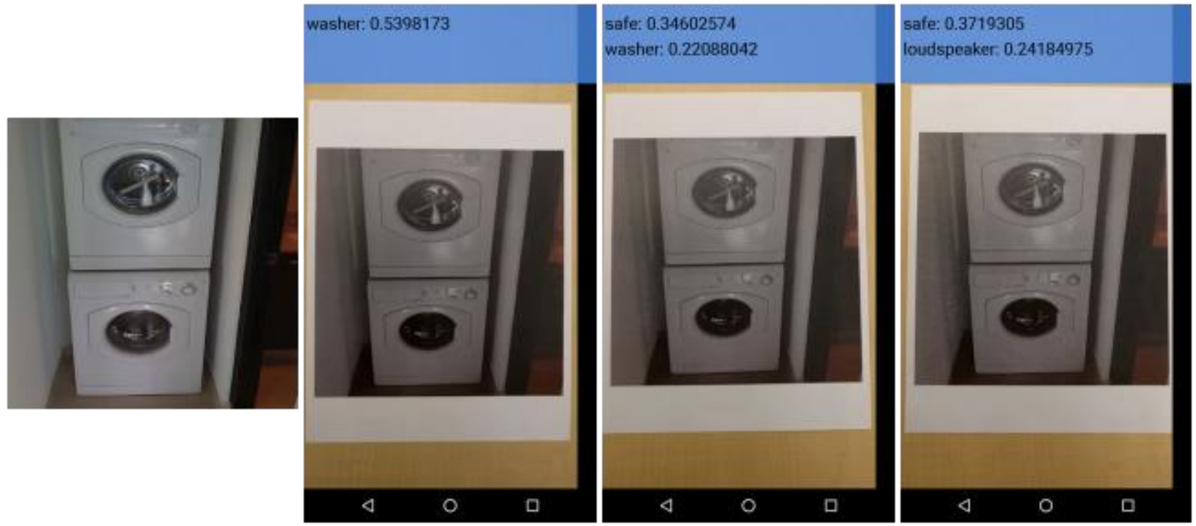

(a) Image from dataset (b) Clean image (c) Adv. image, $\epsilon = 4$ (d) Adv. image, $\epsilon = 8$

https://blog.openai.com/adversarial-example-research/

Bieten eine Lösung:

- Saliency Maps

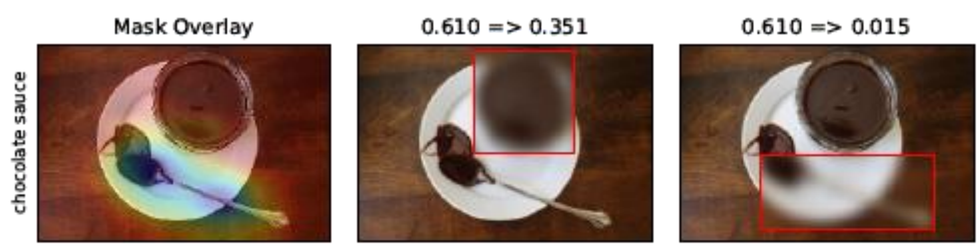

Ruth C. Fong & Andrea Vedaldi, «Interpretable Explanations of Black Boxes by Meaningful Perturbation», 2017

Was? → Wie? → Wo?

3

Wo wird das heute bereits eingesetzt?

Erkennung von Musiknotation
Grundlage für Digitalisierung in Orchestern und Musikschulen

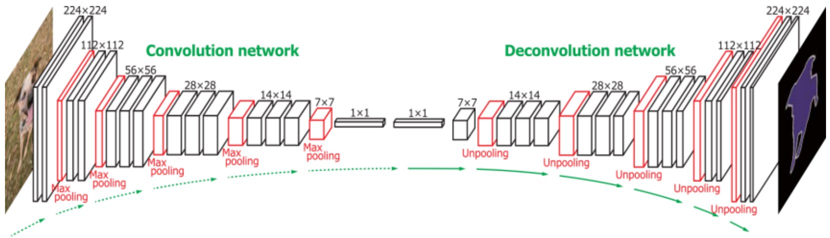

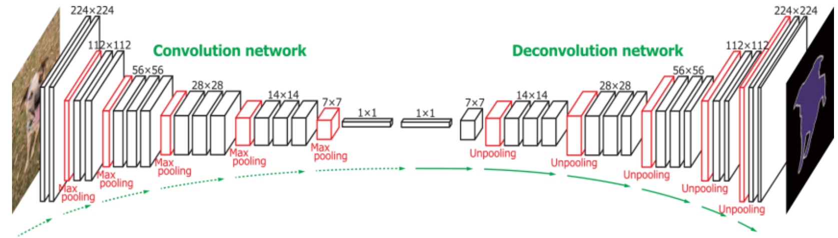

Bild

Label pro Pixel

Objekthaftigkeit

Benannte Symbole

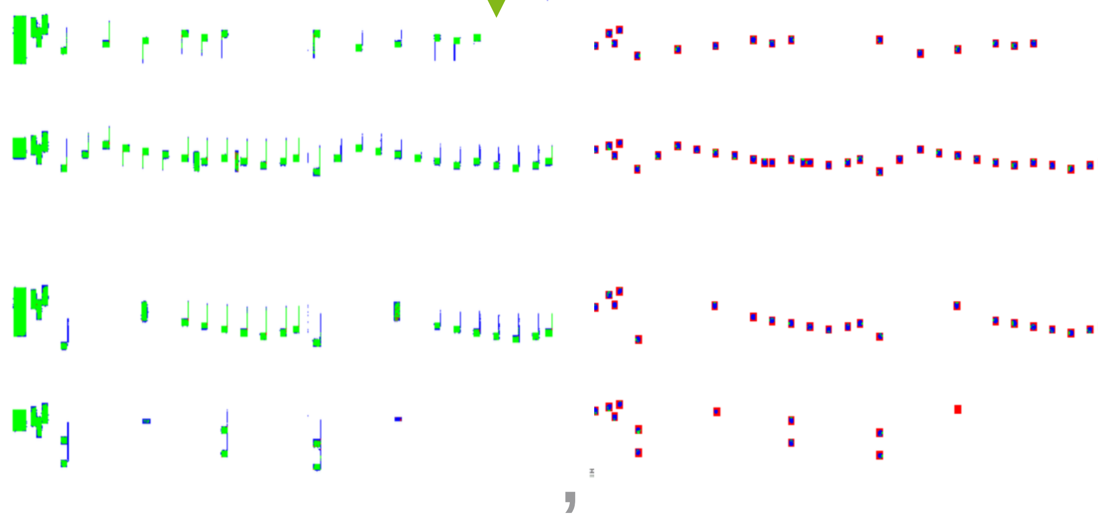

Datengetriebenes Condition Monitoring

Predictive Maintenance von Rotationsmaschinen

Vibrations-Sensor

Merkmalsextraktion

z.B. neuronaler Autoencoder

Früherkennung von Fehlern

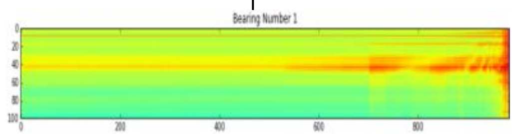

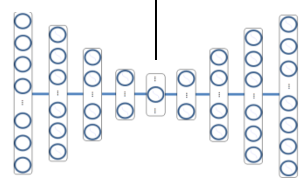

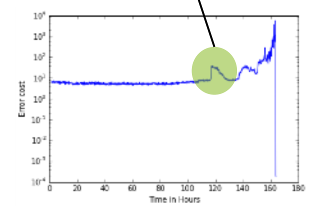

Segmentierung von Zeitungsartikeln

Semiautomatische Medienbeobachtung

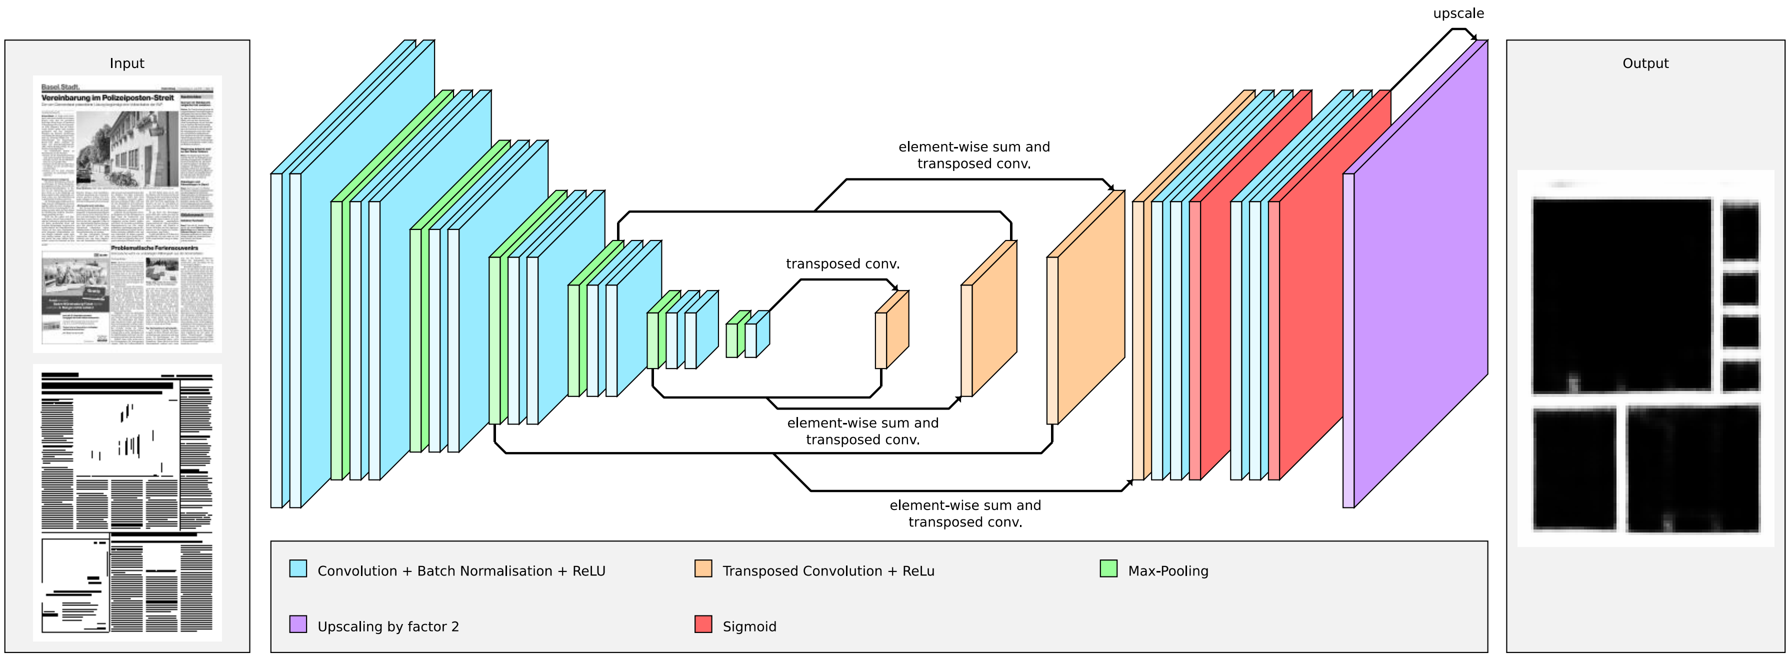

Schlussfolgerungen

- Deep Learning hat zu Paradigmenwechsel in *Mustererkennungsaufgaben* geführt
- Die Zeit vom Grundlagenresultat zur praktischer Anwendung beträgt wenige Monate
- Es gibt Methoden zum Hineinschauen in neuronale Black Boxes
- «~~Denkende~~ rechnende» Maschinen sind trotzdem nur *insel(-hoch-)begabt*

  ➔ *Herausforderungen bestehen im Bereich Robustheit, Interpretierbarkeit, rechtl. Stellung*

datalab
www.zhaw.ch/datalab

swiss group for artificial intelligence and cognitive science
SGAICO

data service
Swiss Alliance for Data-Intensive Services

Mehr zu mir:

- Leiter ZHAW Datalab, Vice President SGAICO, Board Data+Service
- thilo.stadelmann@zhaw.ch
- 058 934 72 08
- www.zhaw.ch/~stdm

Mehr zum Thema:

- KI: https://sgaico.swissinformatics.org/
- Verband Data+Service Science: www.data-service-alliance.ch
- Gemeinsame Projekte: datalab@zhaw.ch

➔ Fragen Sie gerne an.

ANHANG

Künstliche Intelligenz

WaveNet lässt Computersprache natürlich klingen

von Henning Steier / 12.9.2016, 10:05 Uhr

Die Google-Tochter DeepMind hat ein neuronales Netz präsentiert, das Rechner fast wie Menschen klingen lässt. Es macht auch Musik.

KOMMENTARE

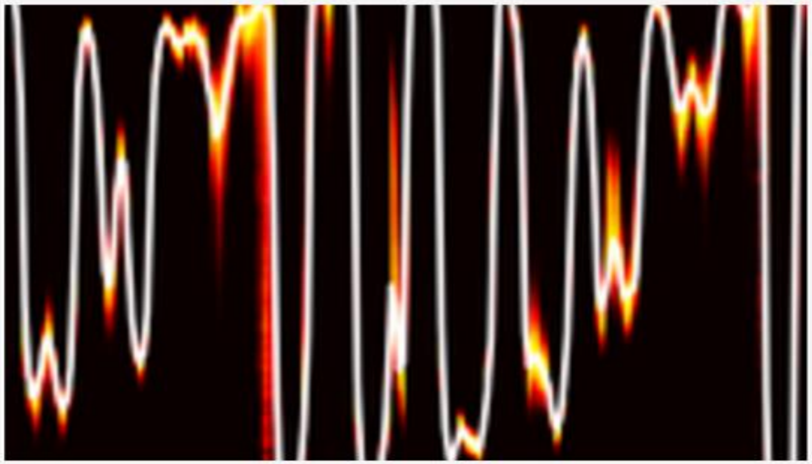

DeepMind lässt WaveNet Sprachwellen erzeugen. (Symbolbild: PD)

Die Google-Tochter DeepMind machte zuletzt mit ihrem Sieg beim Spiel «Go» Schlagzeilen: Ihre Software AlphaGo schlug im Frühjahr einen der besten menschlichen Spieler, Lee Sedol. Nun hat das Londoner Unternehmen WaveNet präsentiert: Dieses neuronale Netz erzeugt Sprache, die sehr natürlich klingt – zumindest wenn man die im Blogeintrag des Unternehmens zu hörenden Klangbeispiele als Massstab nimmt. Man hat sogar das Gefühl, Atempausen zu hören.

MEISTGELESEN

Künstliche Intelligenz
Kein Google für jeden
KOMMENTAR / Henning Steier / 5.10.2016

Neue Produkte aus Mountain View
Google macht sich nicht nur im Wohnzimmer breit
Henning Steier / 4.10.2016

Dropbox
68 Millionen verschlüsselte Passwörter im Netz
5.10.2016

Zürcher Hochschule für Angewandte Wissenschaften

Generierte Sprache «aus Texteingabe»

Generierte Musik «ohne Inhaltsvorgabe»

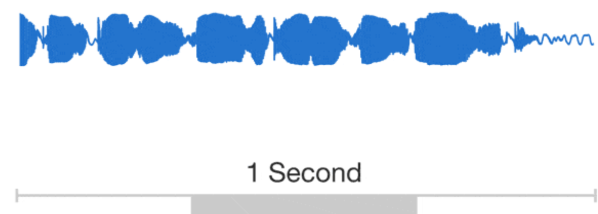

Künstliche Intelligenz

WaveNet lässt Computersprache natürlich klingen

von Henning Steier / 12.9.201

Die Google-Tochter DeepM
macht auch Musik.

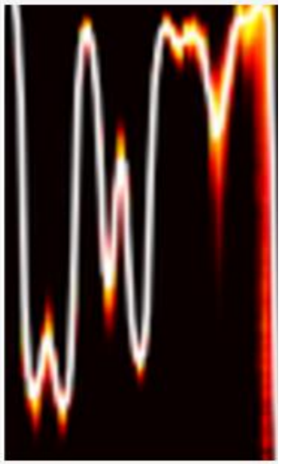

DeepMind lässt WaveNet Spra

Die Google-Tochter Deep
Spiel «Go» Schlagzeilen:
einen der besten mensch
Londoner Unternehmen
erzeugt Sprache, die seh
im Blogeintrag des Unte
Massstab nimmt. Man ha

Intro

What if you could imitate a famous celebrity's voice or sing like a famous singer? This project started with a goal to convert someone's voice to a specific target voice. So called, it's voice style transfer. We worked on this project that aims to convert someone's voice to a famous English actress Kate Winslet's voice. We implemented a deep neural networks to achieve that and more than 2 hours of audio book sentences read by Kate Winslet are used as a dataset.

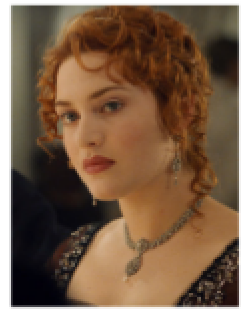

Model Architecture

This is a many-to-one voice conversion system. The main significance of this work is that we could generate a target speaker's utterances without parallel data like <source's wav, target's wav>, <wav, text> or <wav, phone>, but only waveforms of the target speaker. (To make these parallel datasets needs a lot of effort.) All we need in this project is a number of waveforms of the target speaker's utterances and only a small set of <wav, phone> pairs from a number of anonymous speakers.

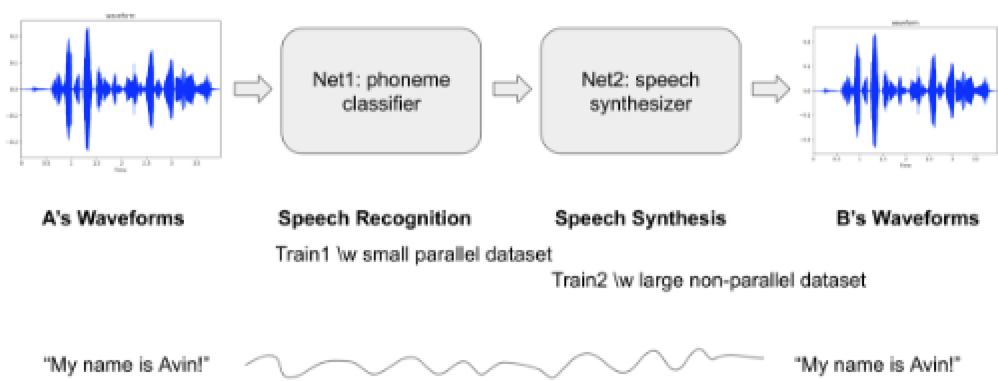

Zürcher Hochschule
für Angewandte Wissenschaften

nerierte Sprache
s Texteingabe»

nerierte Musik
ne Inhaltsvorgabe»

Idee: Mehr Tiefe zum Lernen von Merkmalen

Klassische Bild-verarbeitung

Merkmalsextraktion
(SIFT, SURF, LBP, HOG, etc.)

(0.2, 0.4, …)

(0.4, 0.3, …)

Klassifikation
(SVM, Neuronales Netz, etc.)

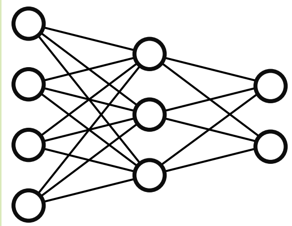

Containerschiff

Tiger

…

Mit Convolutional Neural Networks (CNNs)

Nimmt rohe Pixel entgegen,
Merkmale werden mitgelernt!

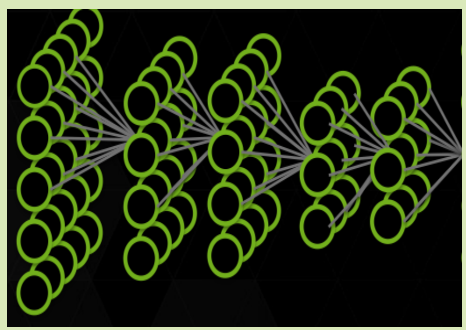

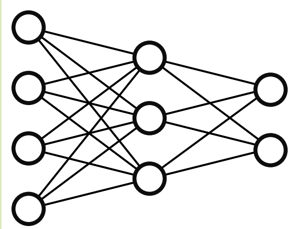

Containerschiff

Tiger

…

Was «sieht» das Neuronale Netz?
Hierarchien komplexer werdender Merkmale

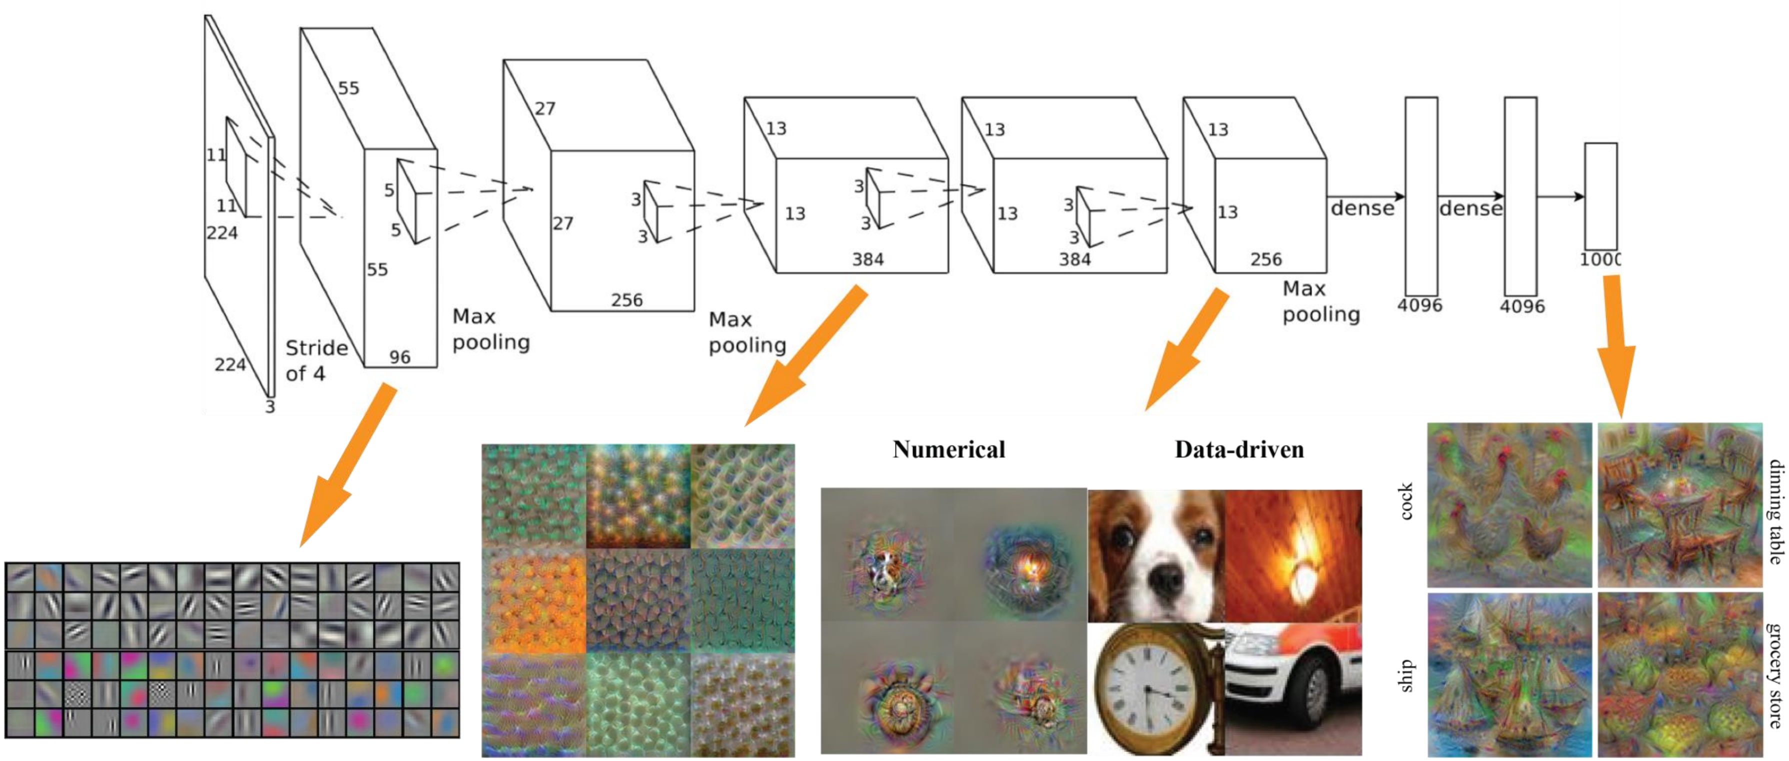

Conv 1: Edge+Blob | Conv 3: Texture | Conv 5: Object Parts | Fc8: Object Classes

Quelle: http://vision03.csail.mit.edu/cnn_art/data/single_layer.png