Zürcher Hochschule
für Angewandte Wissenschaften

Wie maschinelles Lernen den Markt verändert

Forum Christlicher Wirtschaftswissenschaftler, DE-Giessen, 07. März 2019
10. Arbeitstreffen: Geschäftsmodell “Digital Daten”

Thilo Stadelmann

Was→ Wo?→ Wohin?

1

Was ist künstliche Intelligenz & maschinelles Lernen?

Was ist künstliche Intelligenz?

thinking

"The exciting new effort to **make computers think**... machines with minds, in the full and literal sense." (Haugeland, 1985)

"[The automation of] **activities that we associate with human thinking**, activities such as decision-making, problem solving, learning..." (Bellman, 1978)

"**The study of mental faculties through** the use of **computational models**." (Charniak and McDermott, 1985)

"The study of the **computations that** make it possible to **perceive, reason, and act**." (Winston, 1992)

humanly ← standard measured by → rationally

← concerned with →

"The art **of creating machines that perform** functions that require **intelligence** when performed by people." (Kurzweil, 1990)

"The study of **how to make computers do things** at which, at the moment, **people are better**." (Rich and Knight, 1991)

"Computational Intelligence is **the study of the design of intelligent agents**." (Poole et al., 1998)

"AI... is concerned with **intelligent behaviour in artefacts**." (Nilsson, 1998)

acting

Was ist künstliche Intelligenz?

thinking

"The exciting new effort to **make computers think**... machines with minds, in the full and literal sense." (Haugeland, 1985)

"[The automation of] **activities that we associate with human thinking**, activities such as decision-making, problem solving, learning..." (Bellman, 1978)

"The study of mental faculties through the use of **computational models**." (Charniak and McDermott, 1985)

"The study of the **computations that make it possible to perceive, reason, and act**." (Winston, 1992)

humanly ← Standard measured by → rationally

← concerned with →

"The art **of creating machines that perform functions** that require **intelligence** when performed by people." (Kurzweil, 1990)

"The study of **how to make computers do things** at which, at the moment, **people are better**." (Rich and Knight, 1991)

"Computational Intelligence is **the study of the design of intelligent agents**." (Poole et al., 1998)

"AI... is concerned with **intelligent behaviour in artefacts**." (Nilsson, 1998)

acting

The scientific field concerned with complex computer applications

Was gehört zu künstlicher Intelligenz?

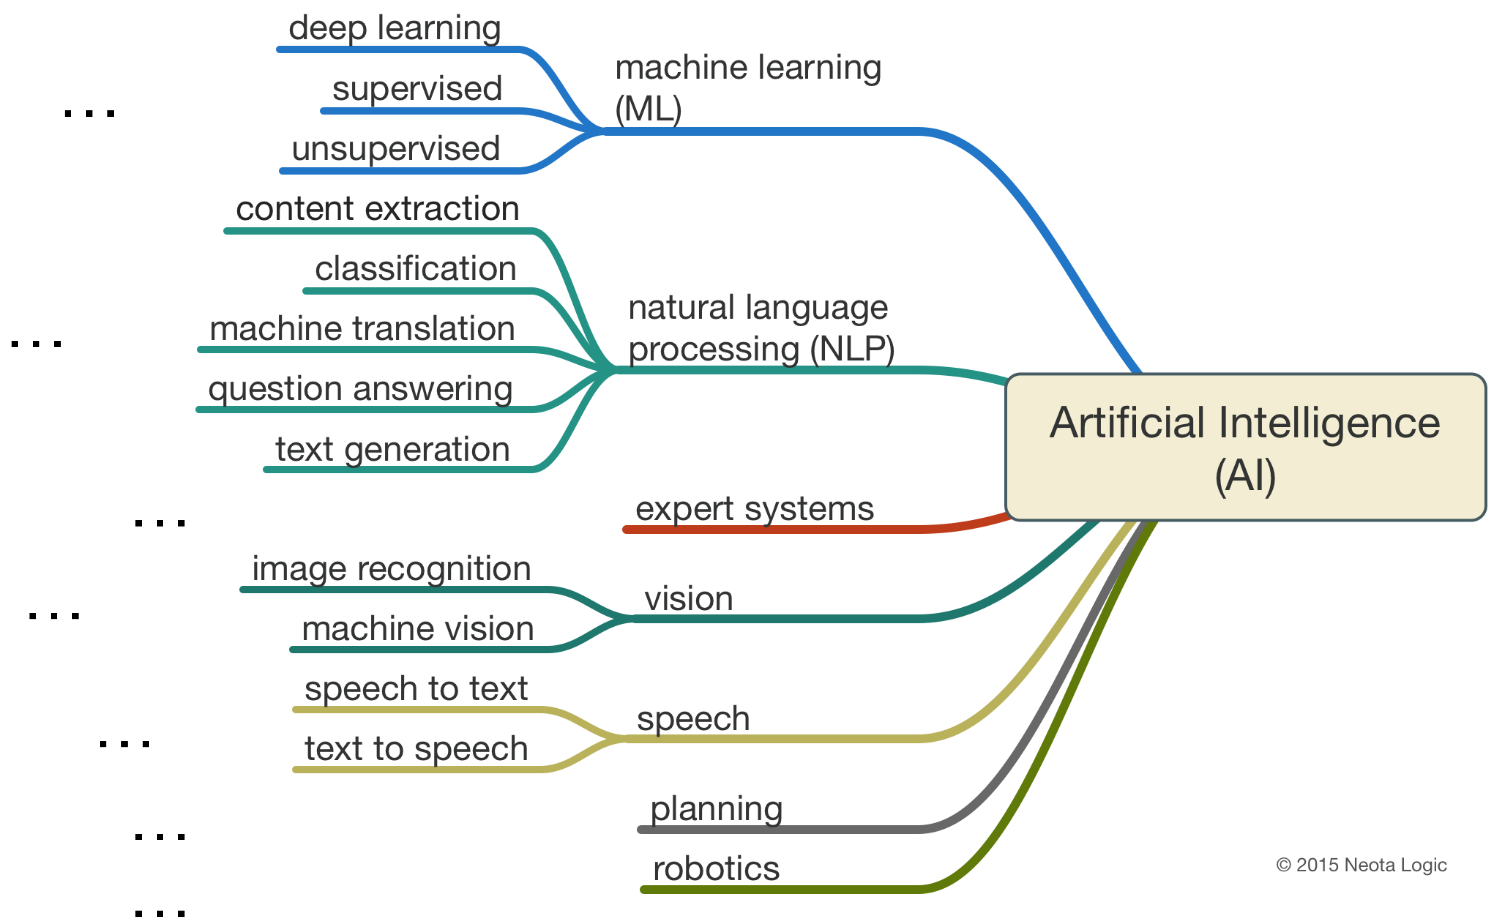

KI im Kontext

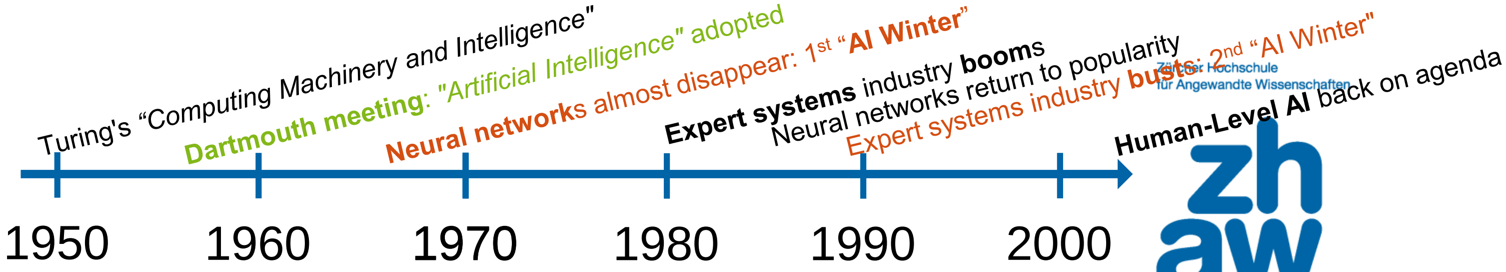

Zürcher Hochschule für Angewandte Wissenschaften

KI im Kontext

- 1950: Turing's "Computing Machinery and Intelligence"
- **Dartmouth meeting**: "Artificial Intelligence" adopted
- **Neural networks** almost disappear: 1st "**AI Winter**"
- **Expert systems** industry **booms**
- Neural networks return to popularity
- Expert systems industry **busts**: 2nd "AI Winter"
- **Human-Level AI** back on agenda

1950 · 1960 · 1970 · 1980 · 1990 · 2000

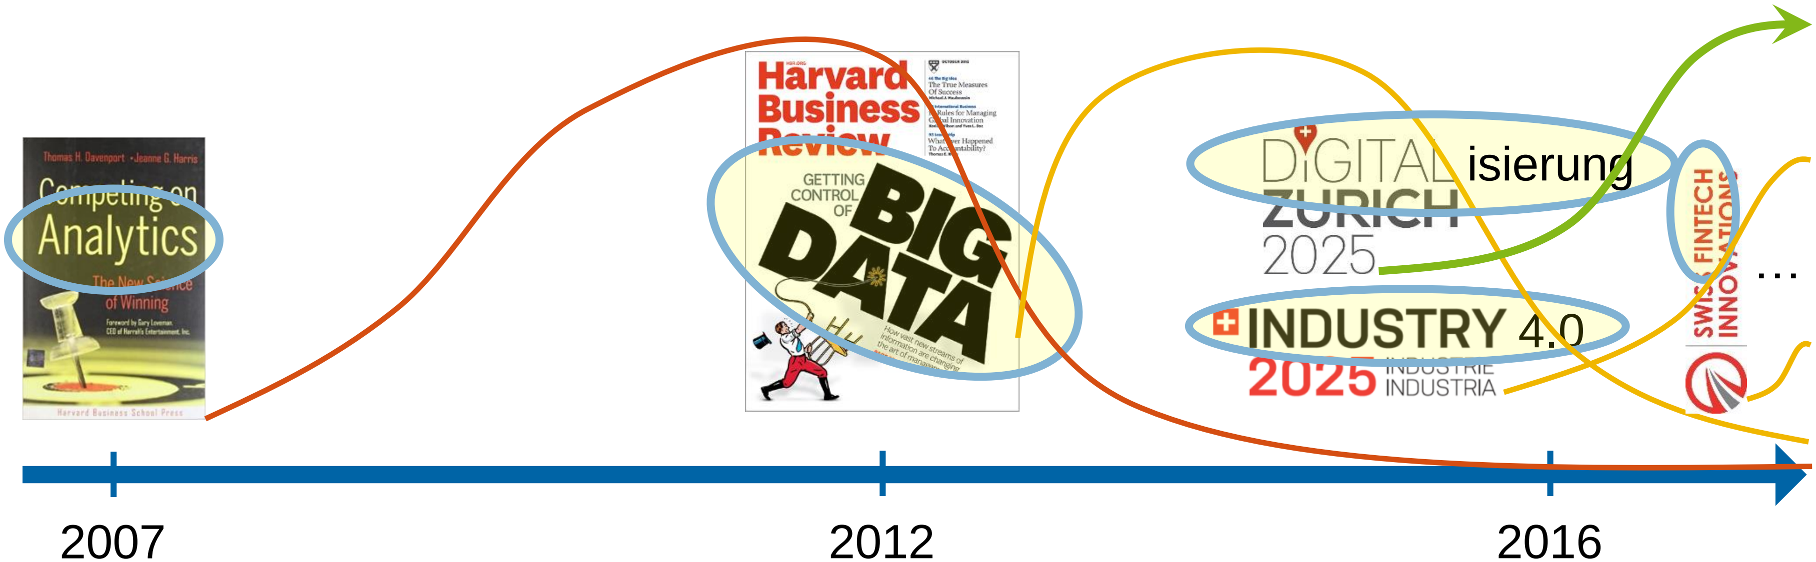

KI im Kontext

Turing's "Computing Machinery and Intelligence"
Dartmouth meeting: "Artificial Intelligence" adopted
Neural networks almost disappear: 1st "**AI Winter**"
Expert systems industry **booms**
Neural networks return to popularity
Expert systems industry **busts**: 2nd "AI Winter"
Human-Level AI back on agenda

1950 1960 1970 1980 1990 2000

Zürcher Hochschule für Angewandte Wissenschaften

Was ist passiert?
Der ImageNet Wettbewerb

1000 Kategorien
1 Mio. Beispiele

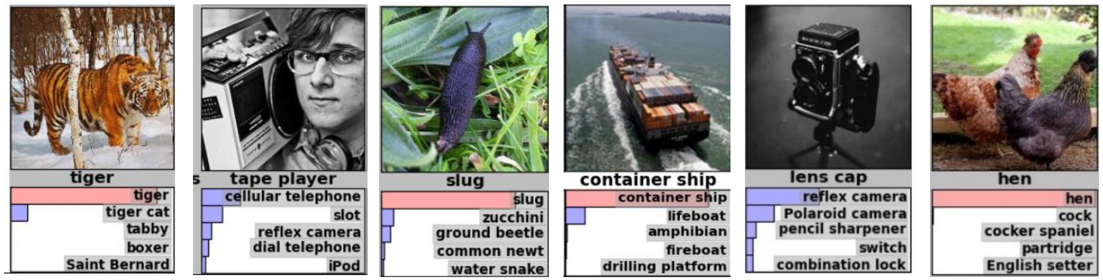

…

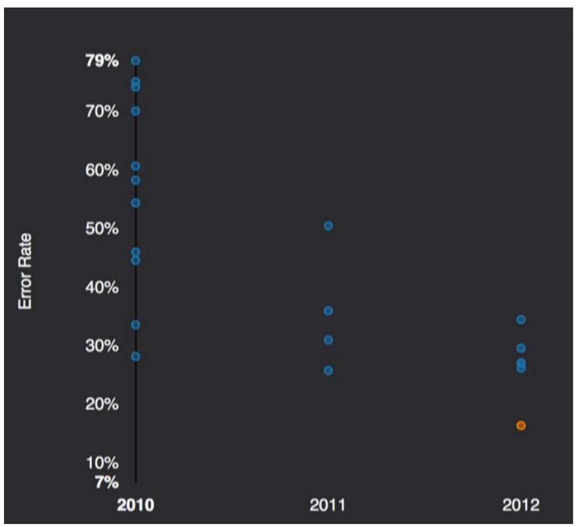

Was ist passiert?
Der ImageNet Wettbewerb

1000 Kategorien
1 Mio. Beispiele

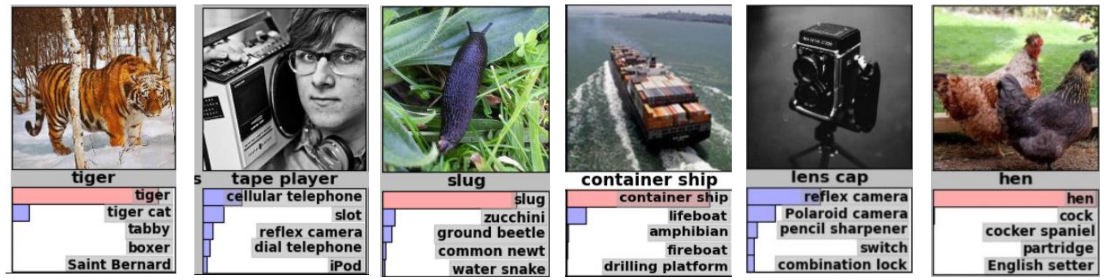

…

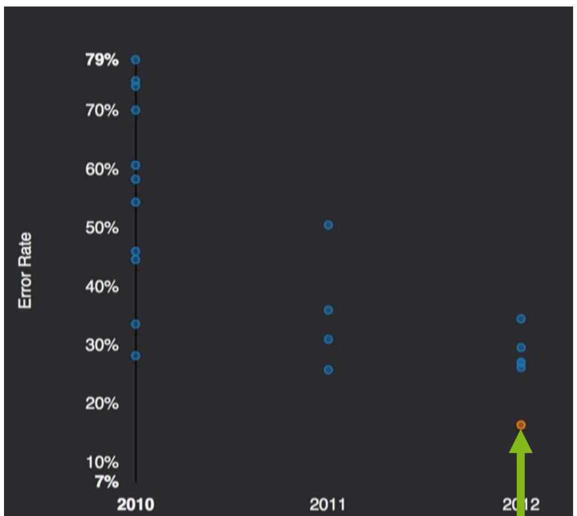

A. Krizhevsky verwendet als erster ein sog. «Deep Neural Network» (CNN)

Was ist passiert?
Der ImageNet Wettbewerb

1000 Kategorien
1 Mio. Beispiele

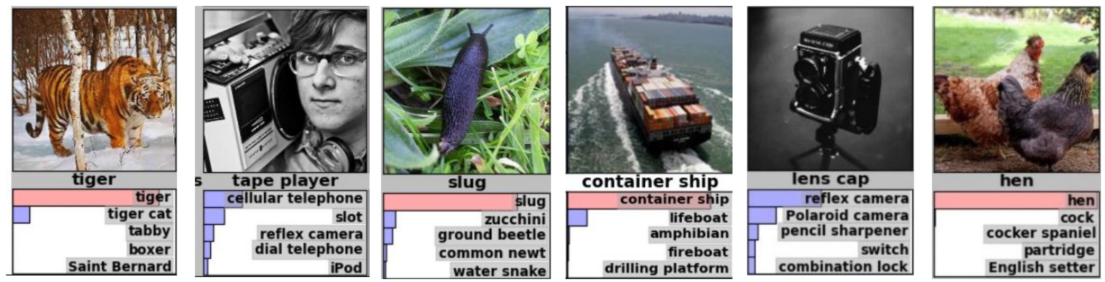

…

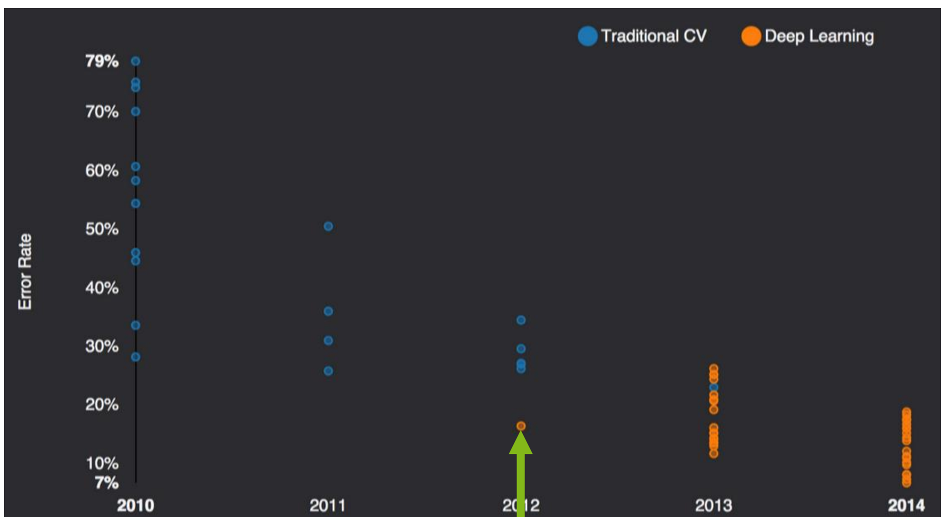

A. Krizhevsky verwendet als erster ein sog. «Deep Neural Network» (CNN)

Was ist passiert?
Der ImageNet Wettbewerb

1000 Kategorien
1 Mio. Beispiele

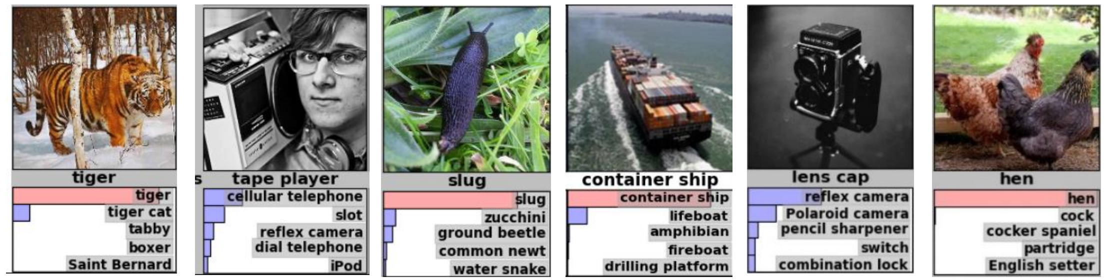

…

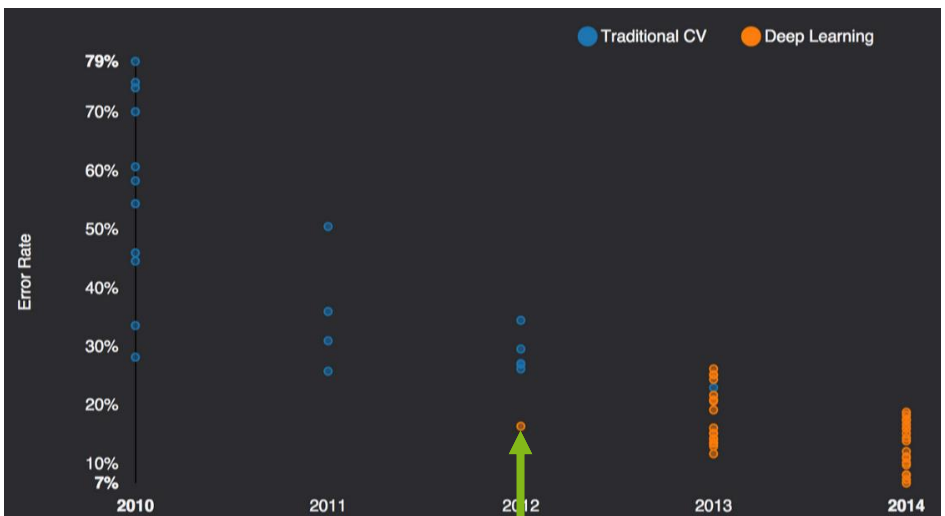

A. Krizhevsky verwendet als erster ein
sog. «Deep Neural Network» (CNN)

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)

Grundidee Deep Learning: “feature learning”

Bildklassifikation (herkömmlicher Ansatz)

Merkmalsextraktion manuell definierter Deskriptoren

(0.2, 0.4, …)

(0.4, 0.3, …)

Traditionelles ML Verfahren (SVM, Neural Network, etc.)

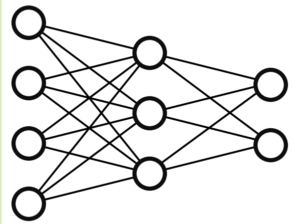

Kontainerschiff

Tiger

Bildclassification (neu: Convolutional Neural Networks)

Rohdaten als Input, wesentliche Merkmale werden automatisch gelernt

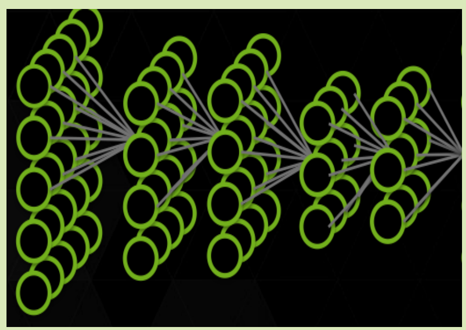

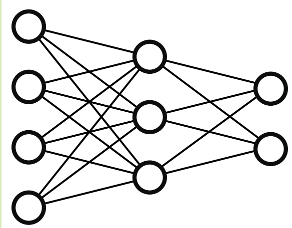

Kontainerschiff

Tiger

Prinzip Machine Learning
“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

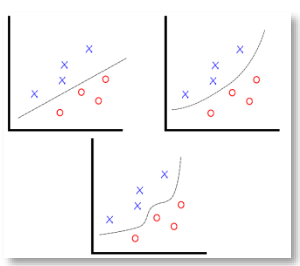

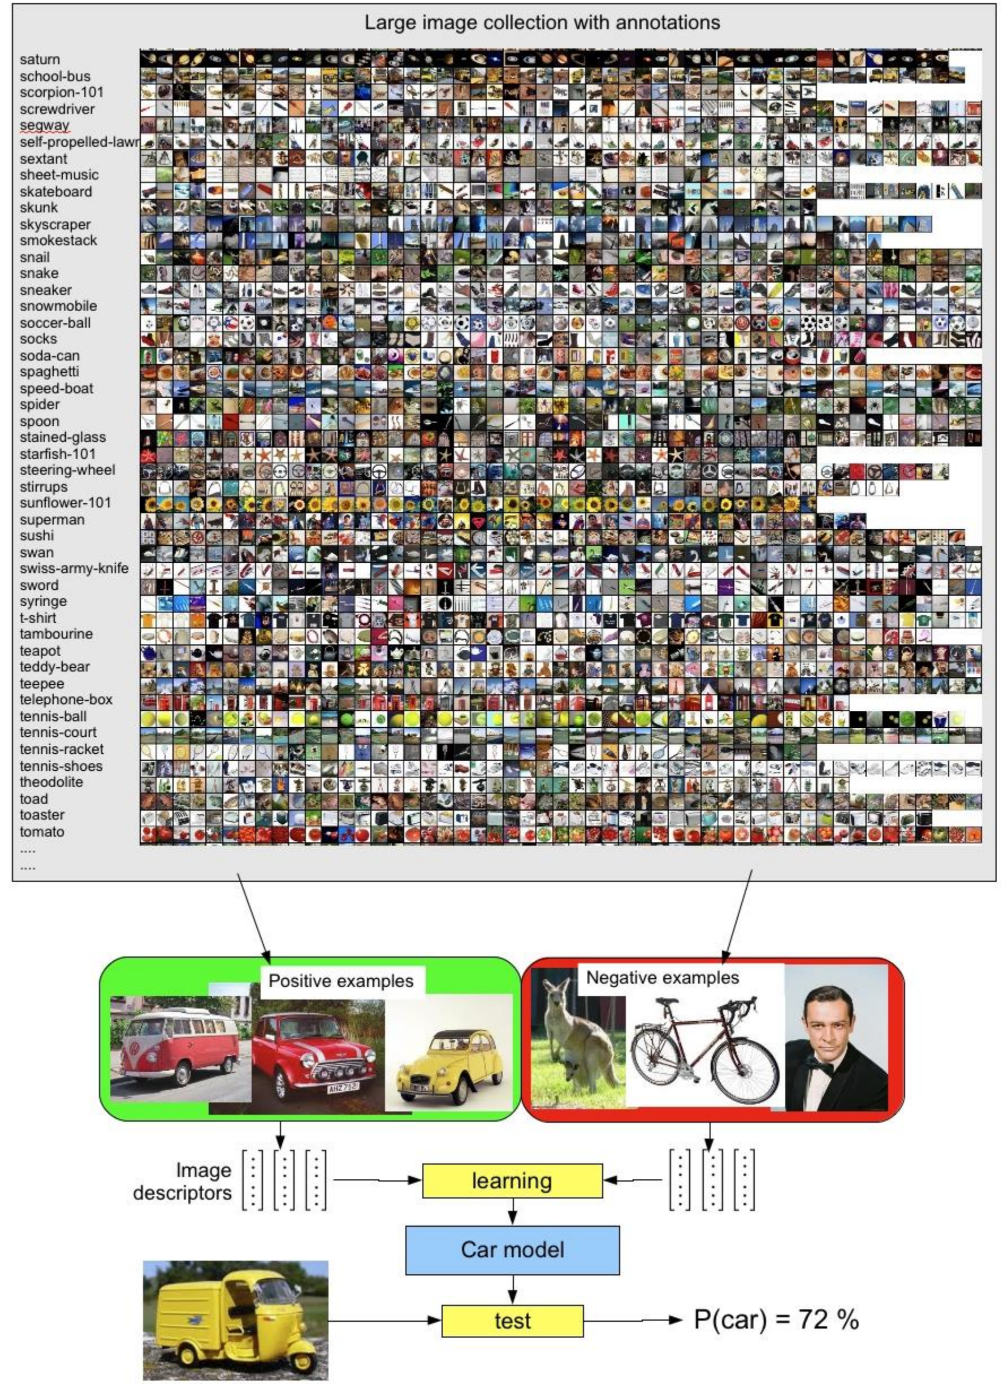

Quelle: http://lear.inrialpes.fr/job/postdoc-large-scale-classif-11-img/attribs_patchwork.jpg

Prinzip Machine Learning
“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

$$f(x) = y$$

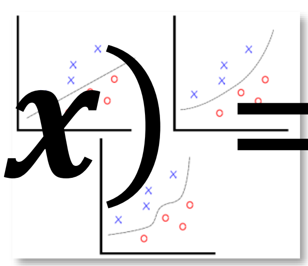

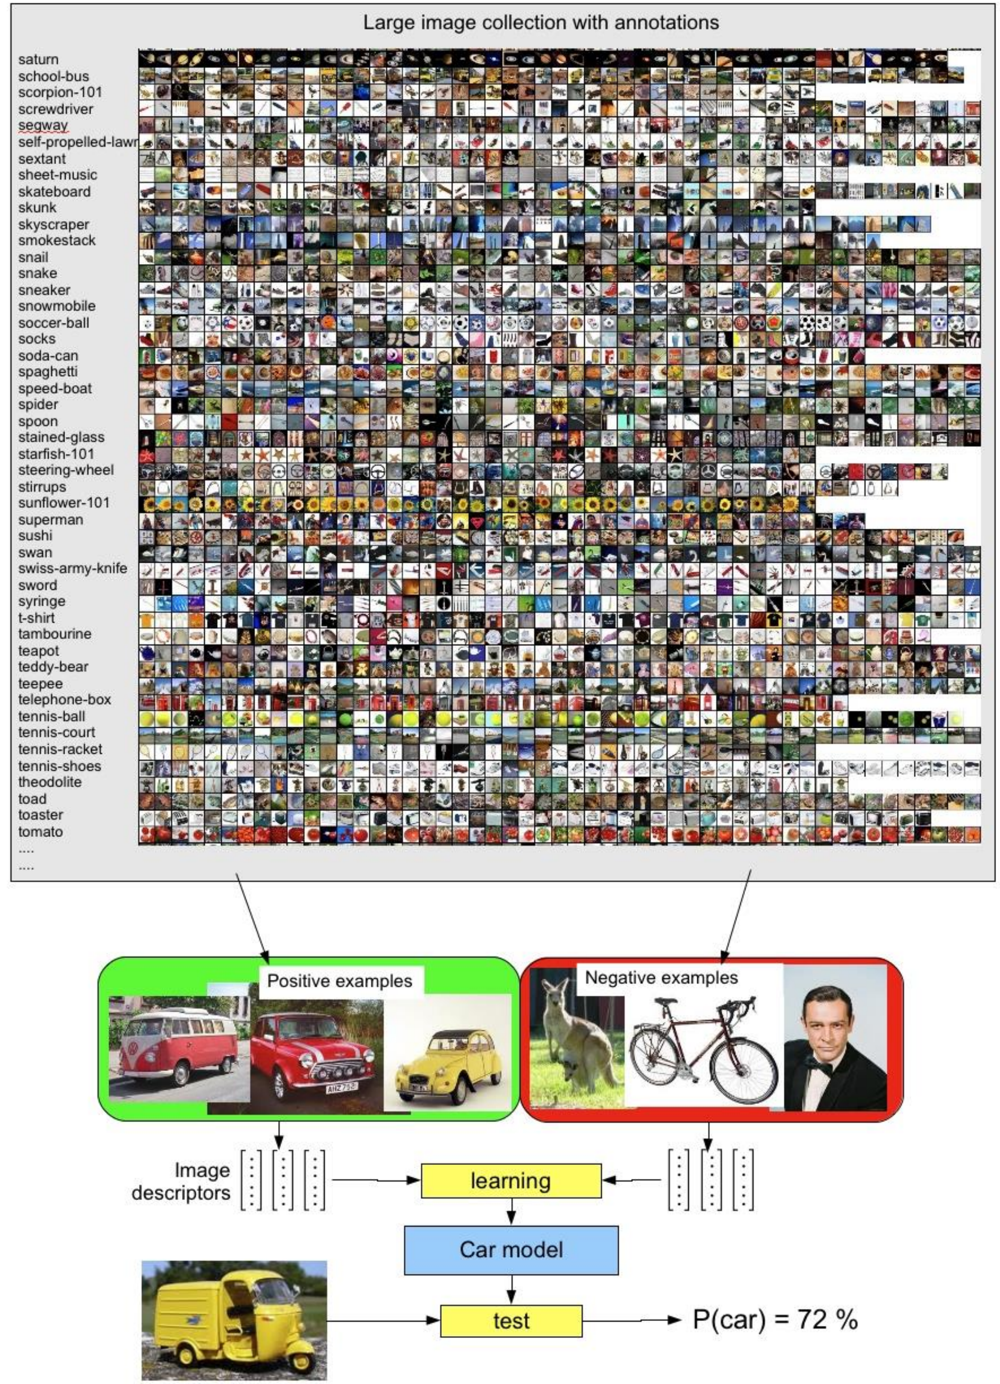

Quelle: http://lear.inrialpes.fr/job/postdoc-large-scale-classif-11-img/attribs_patchwork.jpg

Was→ Wo?→ Wohin?

2

Wo wird das bereits praktisch eingesetzt?

Zürcher Hochschule
für Angewandte Wissenschaften

Einige Beispiele aus den Schlagzeilen


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] I
  - Running DCG
- Step 3: Finding the
  - Image compl
  - [ML-Heavy] I
  - [ML-Heavy] T
  - Completing y
- Conclusion
- Partial bibliography
- Bonus: Incomplete

Introduction

Content-aware fill is a po
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

Zürcher Hochschule
für Angewandte Wissenschaften

Einige Beispiele aus den Schlagzeilen

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

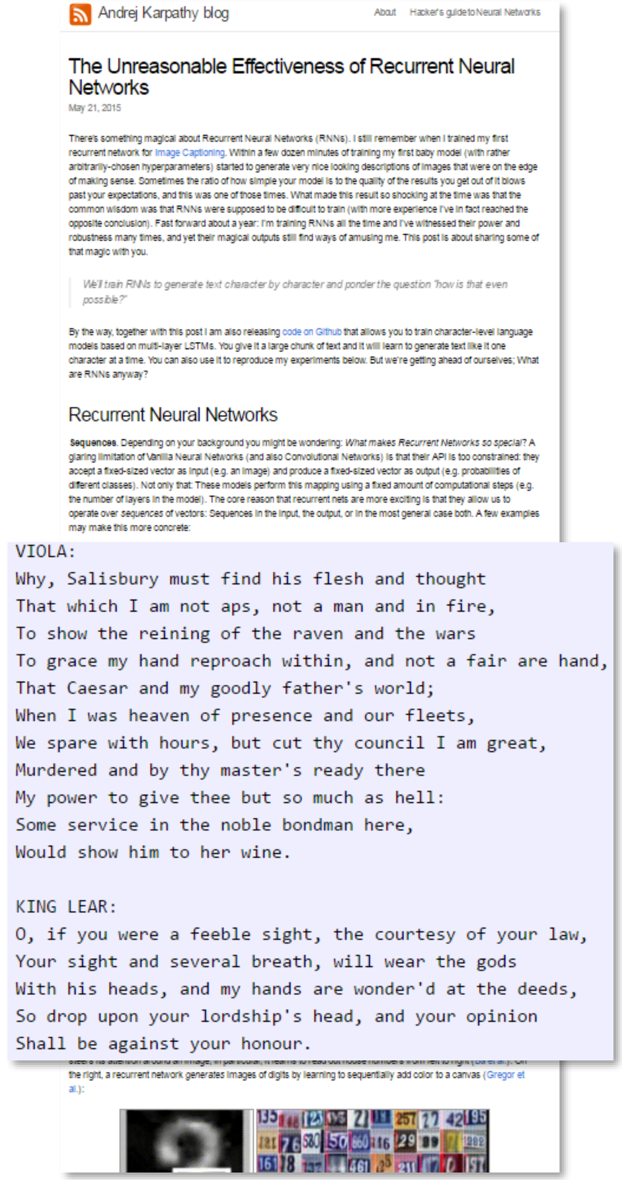

Andrej Karpathy blog About Hacker's guide to Neural Networks

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

Einige Beispiele aus den Schlagzeilen

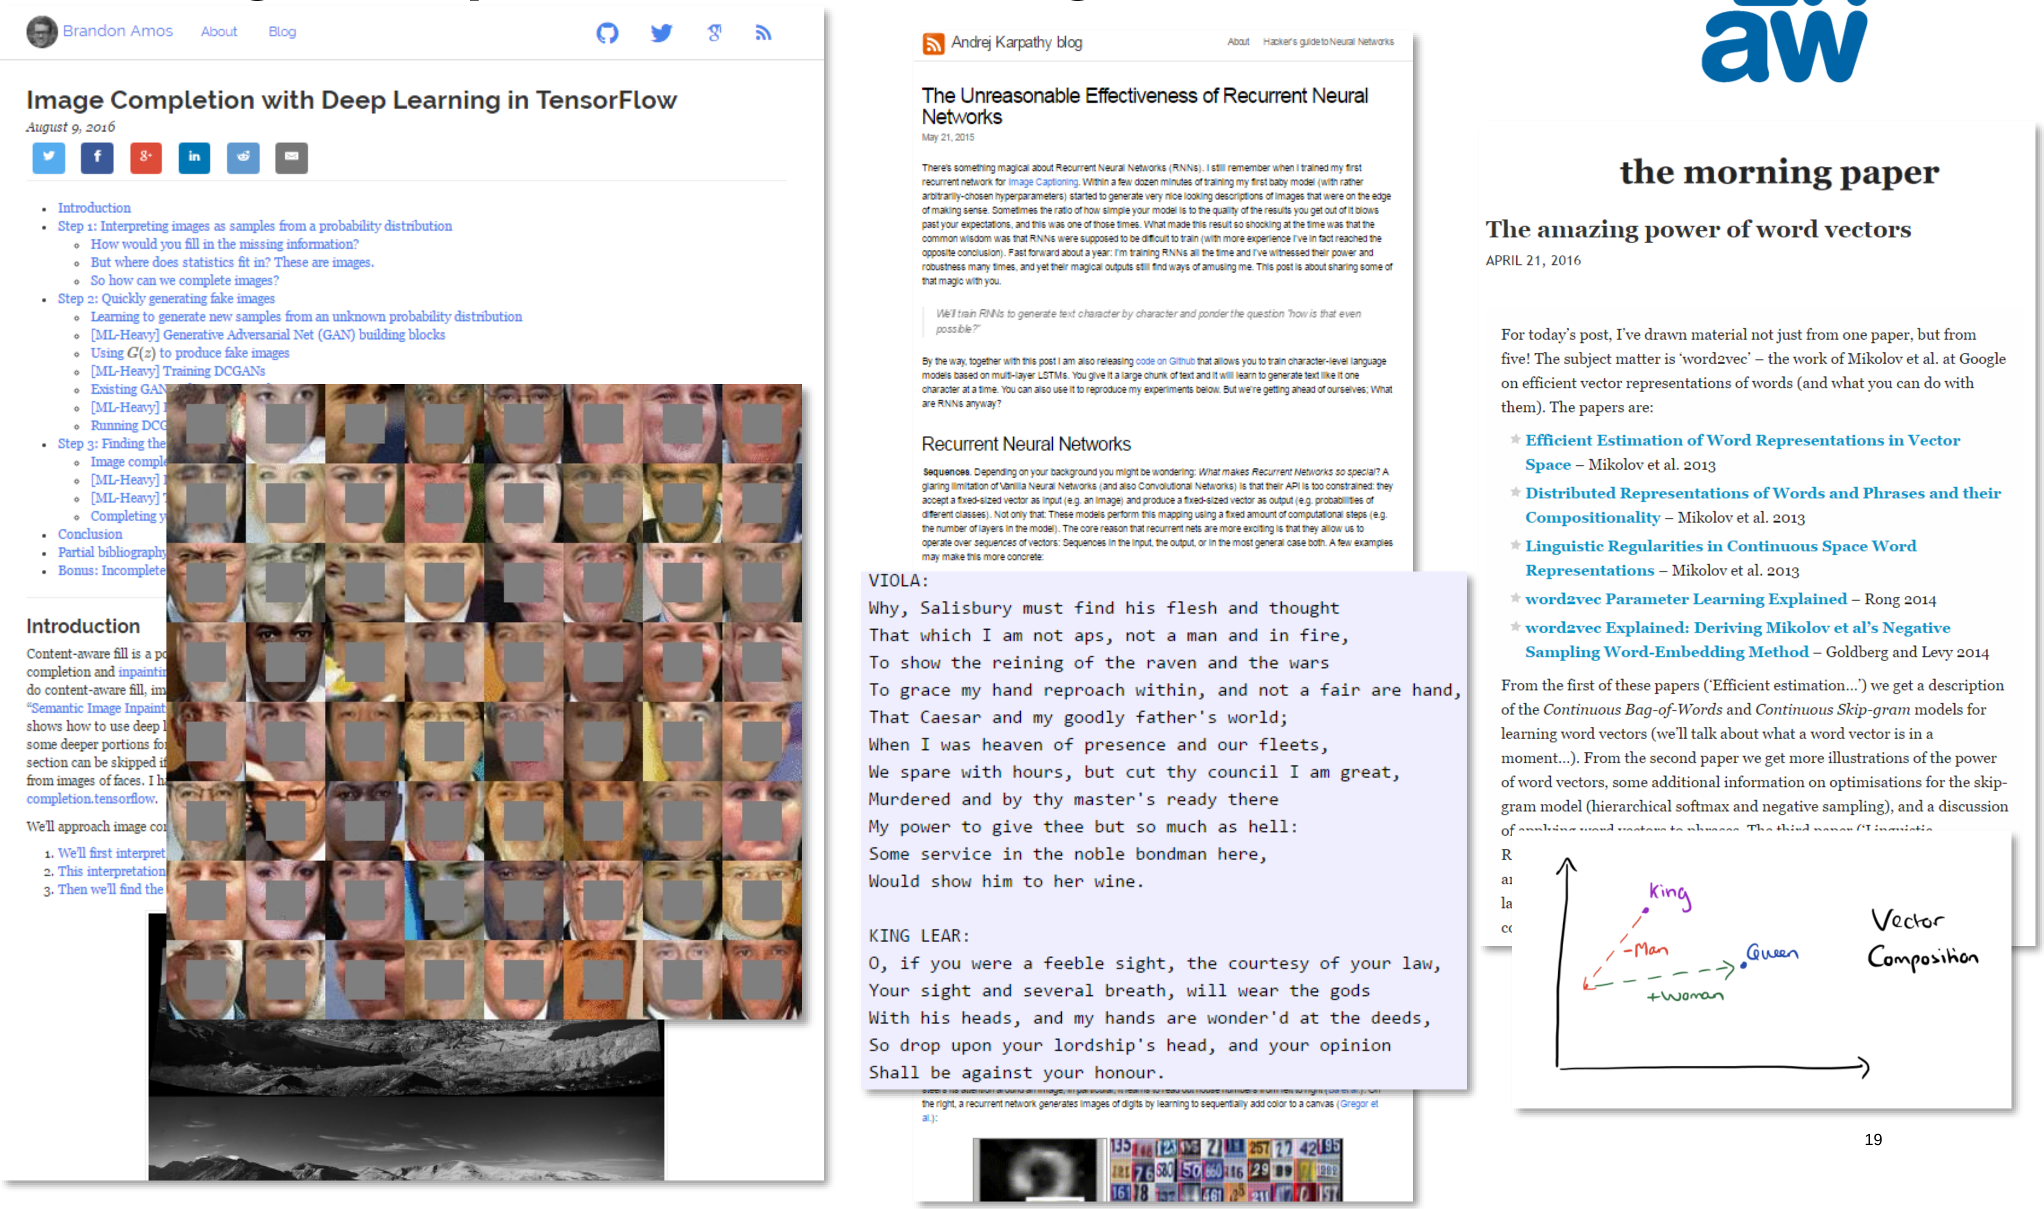

Einige Beispiele aus den Schlagzeilen

Brandon Amos   About   Blog

Image Completion with Deep Learning in TensorFlow

August 9, 2016

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

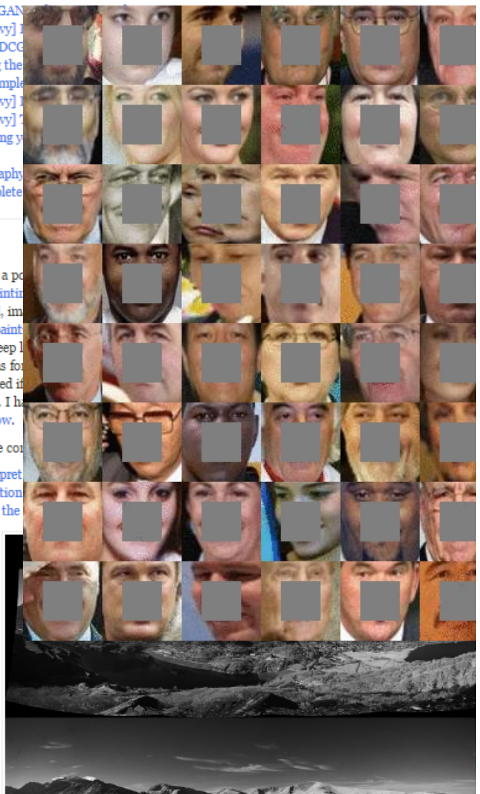

Andrej Karpathy blog   About   Hacker's guide to Neural Networks

The Unreasonable Effectiveness of Recurrent Neural Networks

GEEK.COM

TECH

Nvidia AI Generates Fake Faces Based On Real Celebs

BY STEPHANIE MLOT 10.31.2017 :: 10:00AM EST

32 SHARES

I'm getting a distinctly mid-90s "The Rachel" vibe from the woman in the top left corner (via Nvidia)

STAY ON TARGET

AI Shelley Pens Truly Creepy Horror Stories-And You Can Help

Neural Network Serves Up Truly Frightening Halloween Costume Ideas

Celebrity scandals are about to get a lot more complicated.

Nvidia has **developed** a way of producing photo-quality, AI-generated human profiles—by using famous faces.

hand,

law,

ds,

the morning paper

The amazing power of word vectors

APRIL 21, 2016

For today's post, I've drawn material not just from one paper, but from five! The subject matter is 'word2vec' – the work of Mikolov et al. at Google on efficient vector representations of words (and what you can do with them). The papers are:

- **Efficient Estimation of Word Representations in Vector Space** – Mikolov et al. 2013
- **Distributed Representations of Words and Phrases and their Compositionality** – Mikolov et al. 2013
- **Linguistic Regularities in Continuous Space Word Representations** – Mikolov et al. 2013
- **word2vec Parameter Learning Explained** – Rong 2014
- **word2vec Explained: Deriving Mikolov et al's Negative Sampling Word-Embedding Method** – Goldberg and Levy 2014

From the first of these papers ('Efficient estimation…') we get a description of the *Continuous Bag-of-Words* and *Continuous Skip-gram* models for learning word vectors (we'll talk about what a word vector is in a moment…). From the second paper we get more illustrations of the power of word vectors, some additional information on optimisations for the skip-gram model (hierarchical softmax and negative sampling), and a discussion

R
ar
la
cc

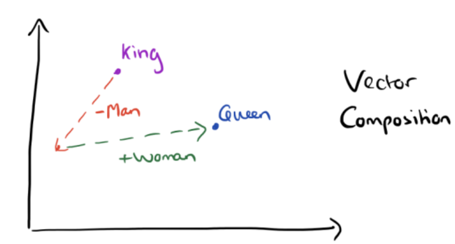

Marktchancen: Machbarkeit vs. Verantwortung

Technologie: Computer Vision mit Deep Learning

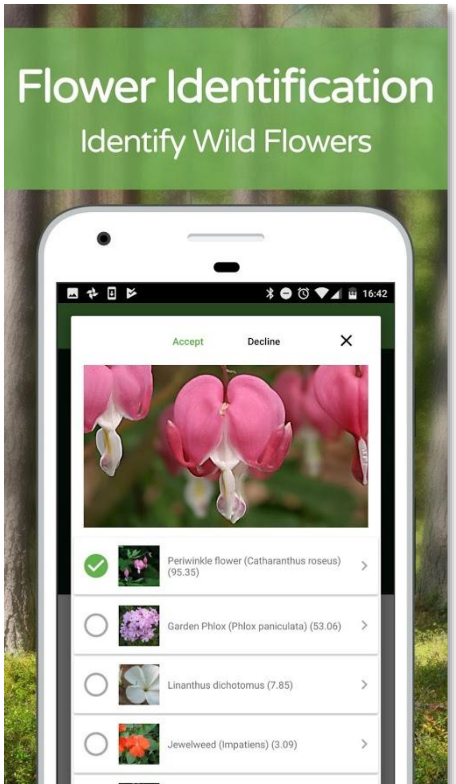

https://www.cultofmac.com/495088/avoid-potentially-deadly-ai-app/

Marktchancen: Machbarkeit vs. Verantwortung
Technologie: Computer Vision mit Deep Learning

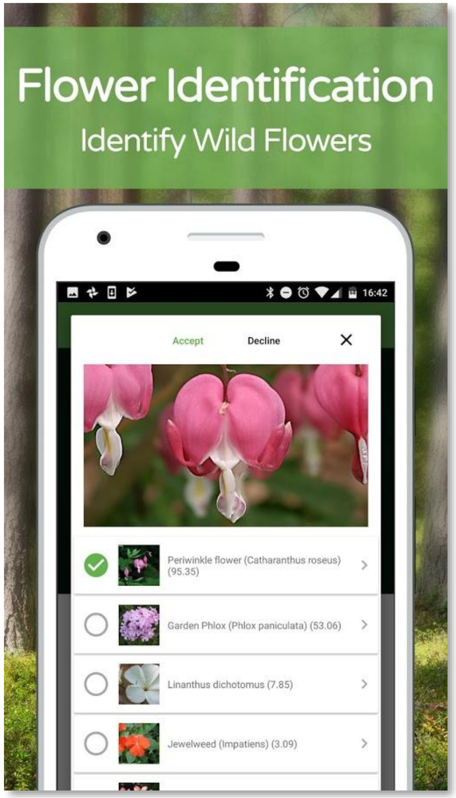

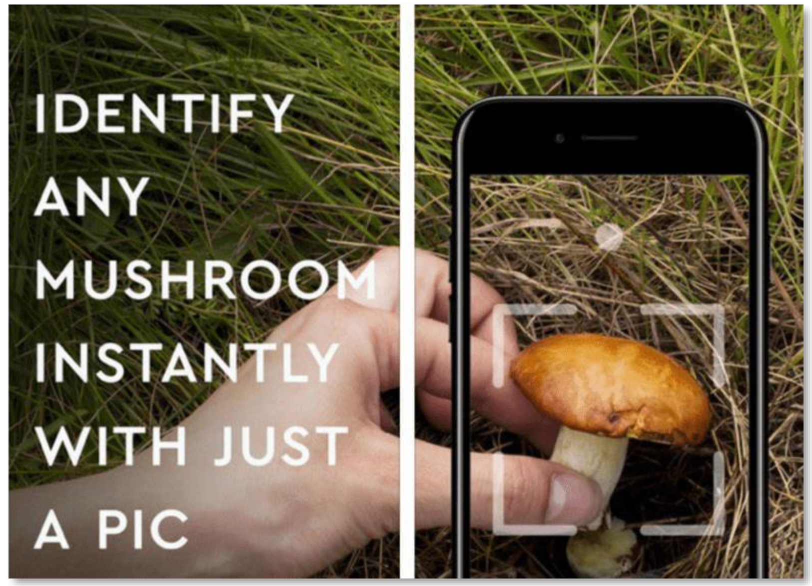

https://www.cultofmac.com/495088/avoid-potentially-deadly-ai-app/

Marktchancen: Markterfolg vs. Regulierung
Technologie: Recommender System

Marktchancen: Markterfolg vs. Regulierung
Technologie: Recommender System

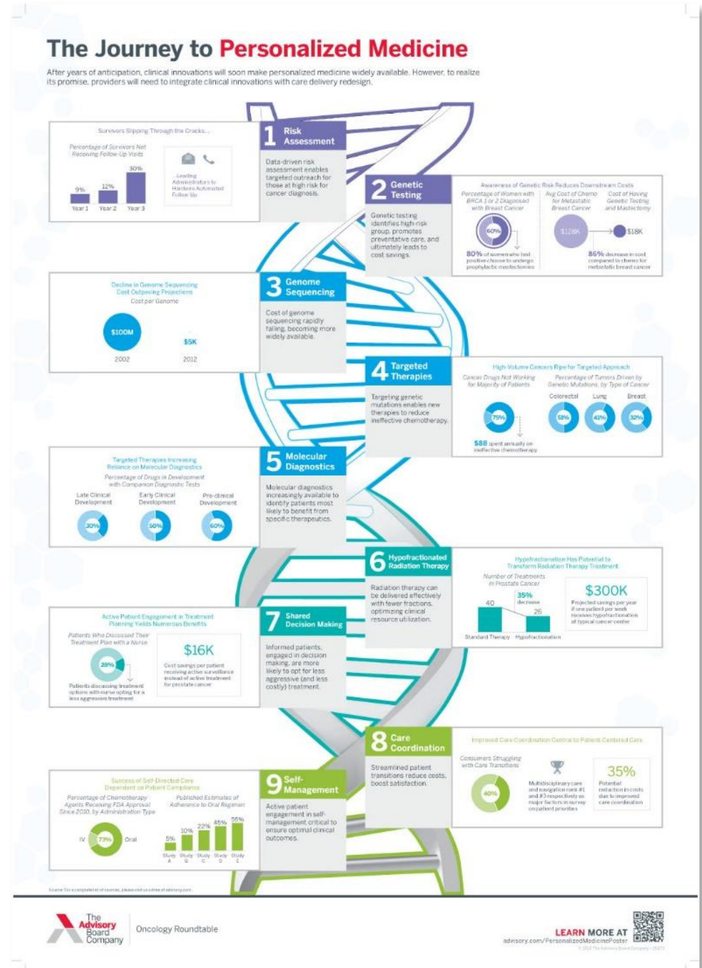

Zurich University of Applied Sciences

Beispiele aus der angewandten Forschung
…mit lokalen Industriepartnern (KMUs)

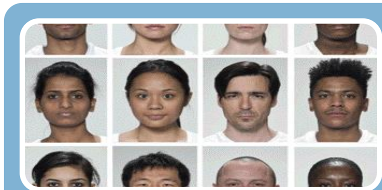

Gesichtserkennung für Stadionzutritt

[!] DEEPIMPACT

- Chance: Enormer Fortschritt in den letzten Jahren
- Herausforderung: Anti-spoofing, algorithmic bias

Automatische Artikelsegmentierung

ARGUS DATA INSIGHTS

- Chance: bildbasiert Layoutregeln lernen
- Herausforderung: Produktisierung, Deployment

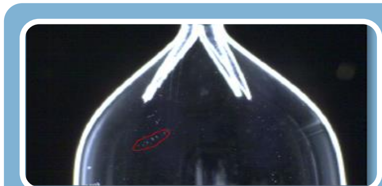

Visuelle Qualitätskontrolle in Produktion

BW-TEC

- Chance: Geschwindigkeit & Präzision
- Herausforderung: hohe Varianz auf „Goldstandart“

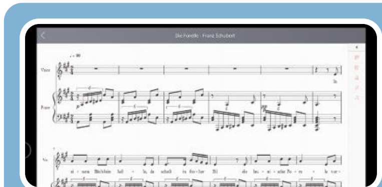

Digitalisierung von Musikalien

SCOREPAD

- Chance: Fortschritt in Digitalisierung von Textdokumenten (OCR)
- Herausforderung: viele kleine Objekte, Kontextabhängigkeit

Grundlagen des disruptiven Potentials (I): Automatisierung "at Scale"

KI

Enorm erweiterte Automatisierungstiefe durch Fortschritt in Mustererkennung

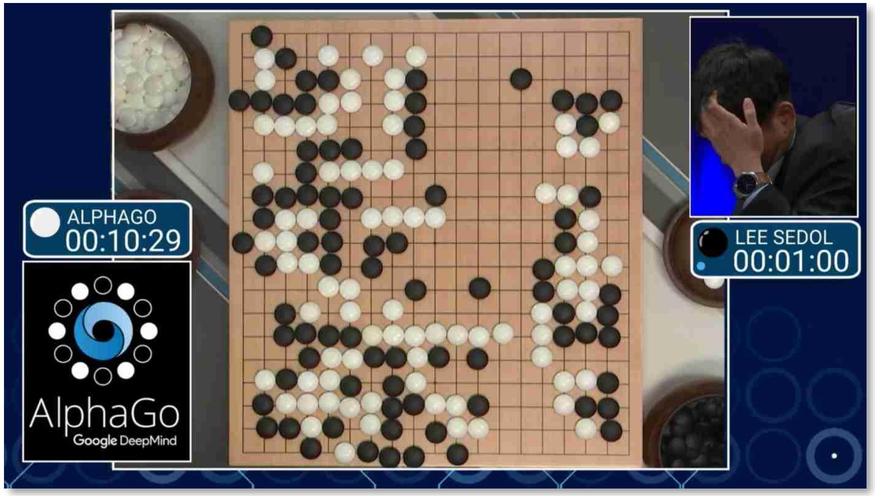

CLOUD COMPUTING

Keine Notwendigkeit mehr für grosse Investitionen in (IT-)Infrastruktur, um in den Markt einzusteigen

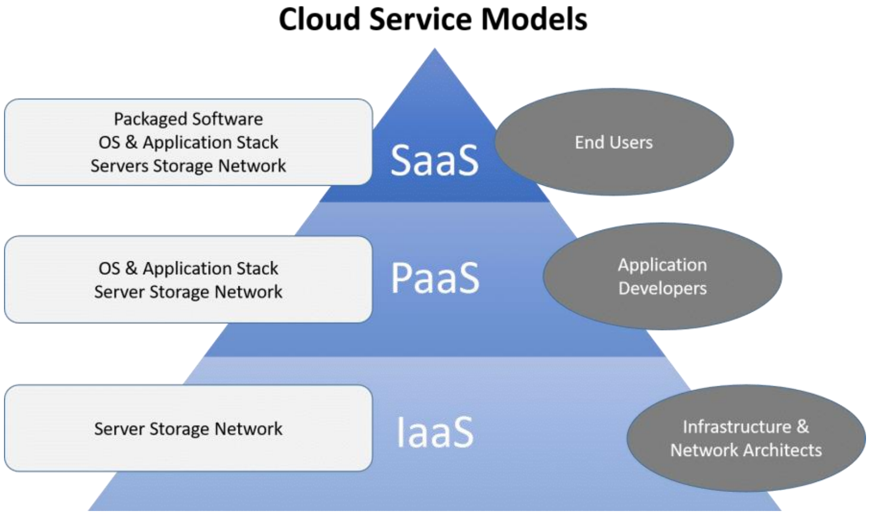

Grundlagen des disruptiven Potentials (II): Entkopplung

Grösse der Idee ≠ Grösse des Unternehmens

…KMUs können **bauen was auch immer sie mögen**
(gegeben Know-how und einen interessanten Business Case)

Technologie ist branchenunabhängig

…was **neue** Kooperationen und Allianzen ermöglicht

Zurich University
of Applied Sciences

Grundlagen des disruptiven Potentials (III): Geschwindigkeit

Durchschnittliche Zeit von Publikation bis Anwendung im Projekt: ca. 3 Monate

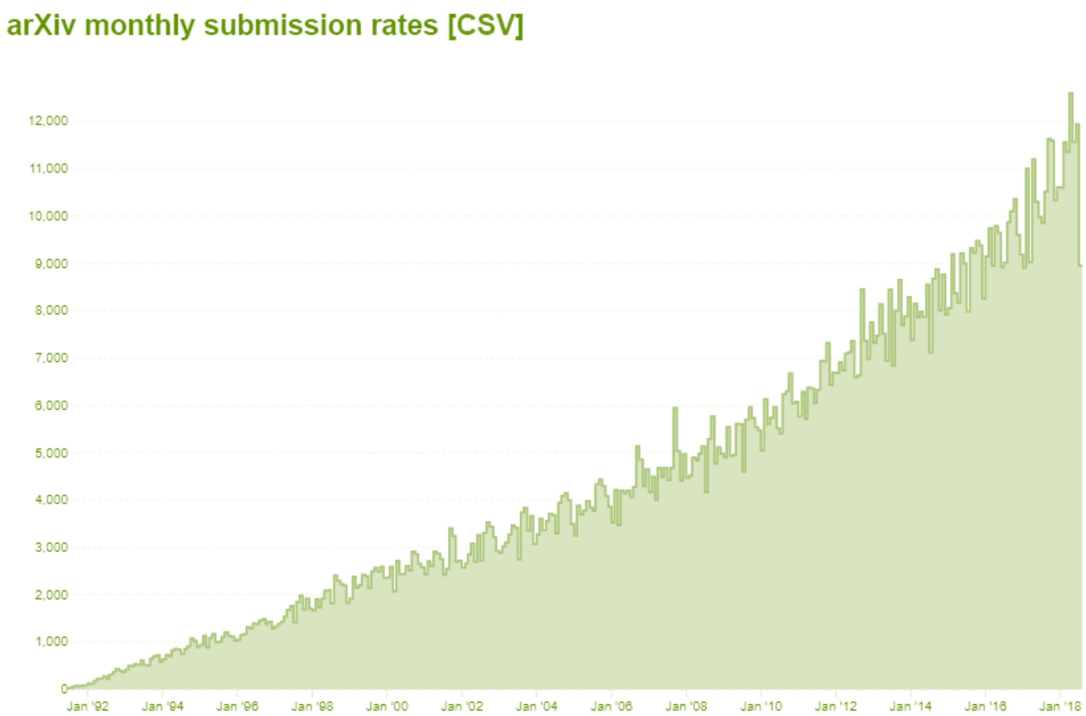

Grundlagen des disruptiven Potentials (III): Geschwindigkeit

Durchschnittliche Zeit von Publikation bis Anwendung im Projekt: ca. 3 Monate

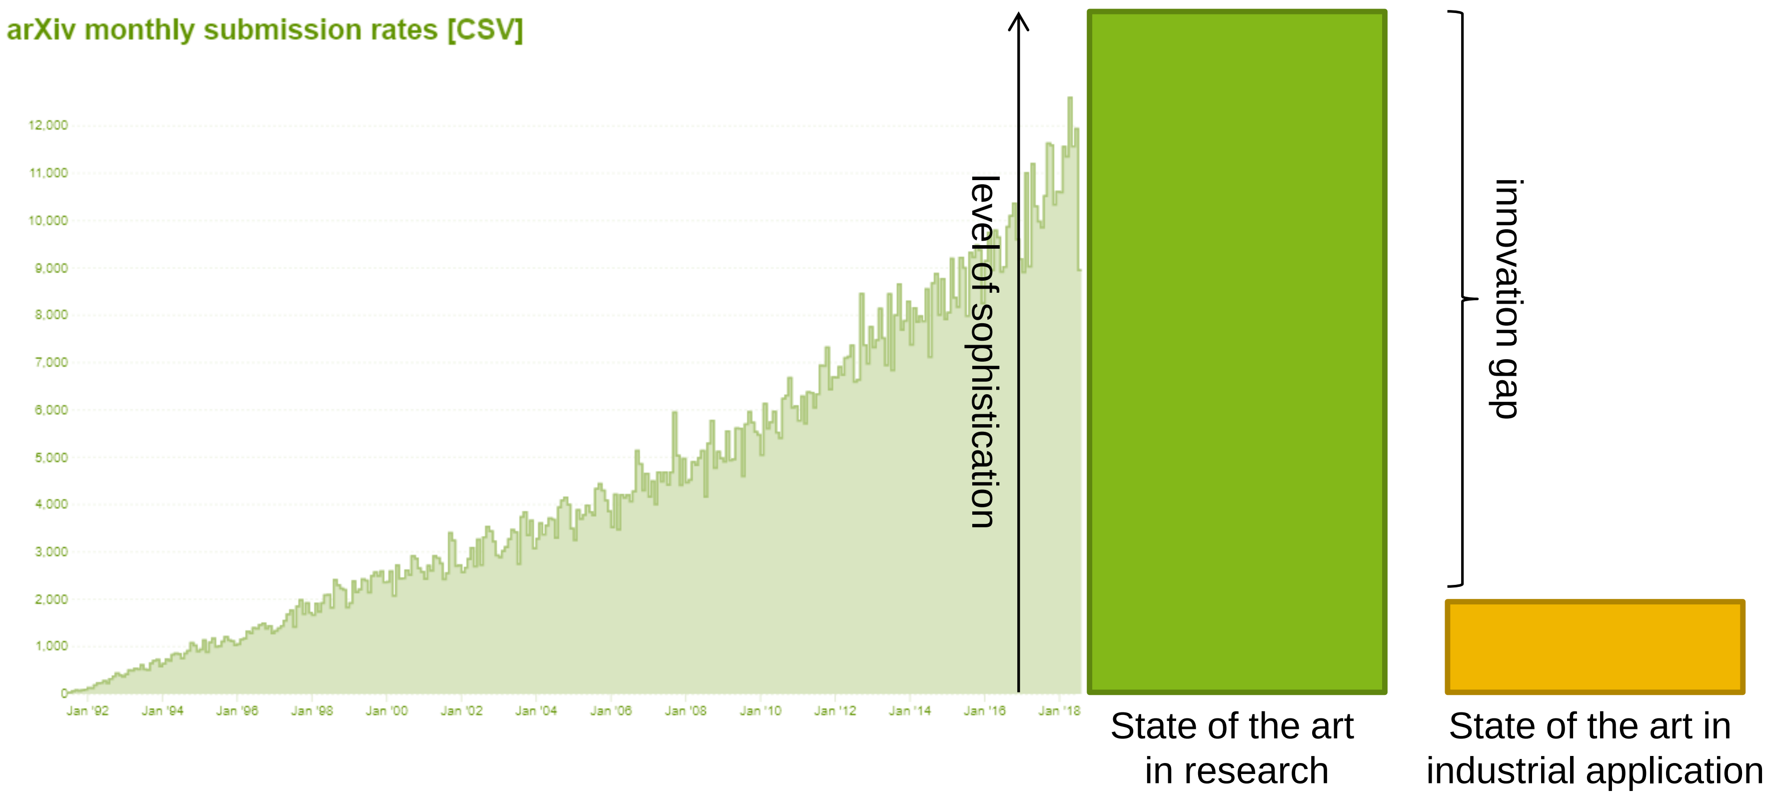

Was→ Wo?→ Wohin?

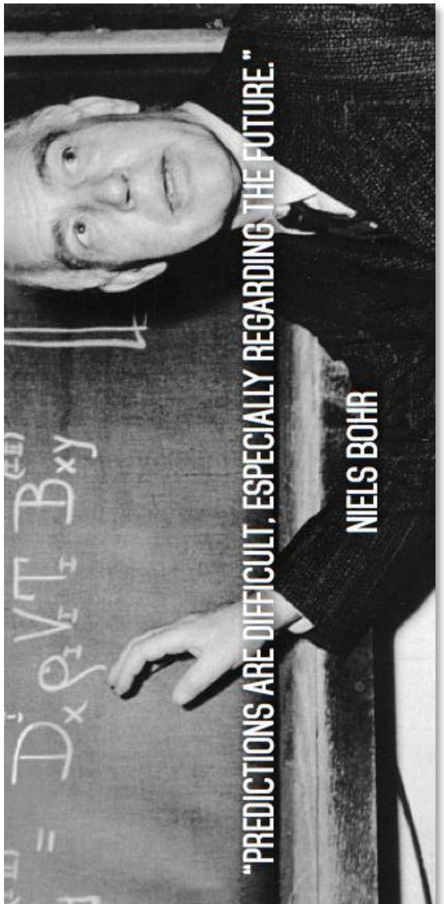

3

Wohin mag das führen?

Aussicht: Disruption
…selbst bei völliger Stagnation des technischen Fortschritts

1. Hypothese: Einsatz (aktueller) KI wird sich massiv ausbreiten (Zeitrahmen: 5 jahre)
- Indikator: **KI-Fortschrit**t momentan hauptsächlich **Industriegetrieben** (**Gewinnaussicht**); Konsumenten kaufen “bequem”; diese Incentivierung “hält den Motor am Laufen”

2. Hypothese: Dies wird unsere Gesellschaften umwälzen
- Kernfragen: Wie **verteilt** sich der algorithmisch (hauptsächlich bei Grosskonzernen) erwirtschaftete **Gewinn**? Wie verteilt sich neue **Freizeit** und **Alltagserleichterung**?

3. Hypothese: Grösste Frage wird der Umgang miteinander sein (nicht der Umgang mit KI)
- Argument: KI (etc.) “for the common good” ist ein wichtiges Thema; entscheidend wird jedoch sein, wie wir **als Gesellschaften die Regeln** für das digitalisierte Zusammenleben (s.o.) **gestalten**



Siehe auch: Stockinger, Braschler & Stadelmann. “Lessons Learned from Challenging Data Science Case Studies”. In: Braschler et al. (Eds), *“Applied Data Science - Lessons Learned for the Data-Driven Business”*, Springer, 2019 (to appear).

Die Vision von Ray Kurzweil
Google, Inc.

The **singularity** is near

- Superintelligence will enhance human life

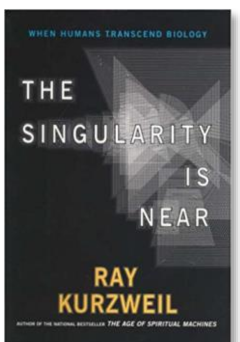

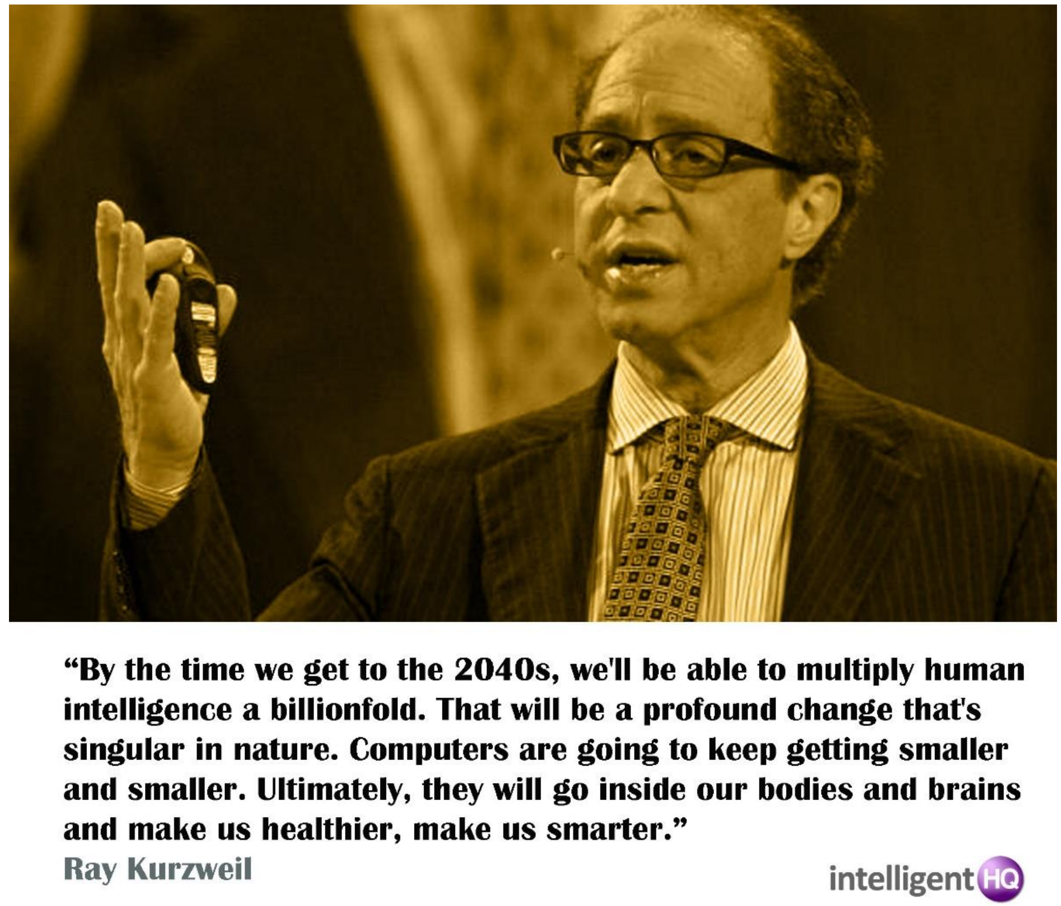

Die Vision von Yuval Noah Harari
Hebrew University of Jerusalem

Humans can become **godlike**

- Humans will upgrade themselves in 3 ways: **biological engineering**, **cyborg** engineering and **robots**

- A new class of people will emerge by 2050: the **useless class** (not just unemployed, but unemployable)
- The most important skill in life will be **learning to learn**: reinvent yourself, again and again until you die to stay out of the useless class
- Computers **function very differently from humans**, and it seems unlikely that computers will become human-like any time soon; however, **intelligence is decoupling from consciousness**

- AI and biotechnology lead to **most powerful narratives** that enable humans to collaborate more effectively and actually **change reality**

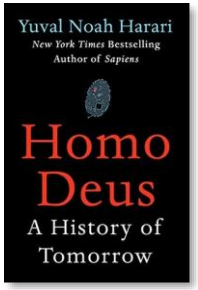

Die Vision von Gene Roddenberry
Creator of Star Trek

„The acquisition of wealth is no longer a driving force in our lives. We ***work to better ourselves and the rest*** *of humanity.“*

Captain Jean-Luc Picard

Compare Richard David Precht‘s *Jäger, Hirten, Kritiker: Eine Utopie für die digitale Gesellschaft.*

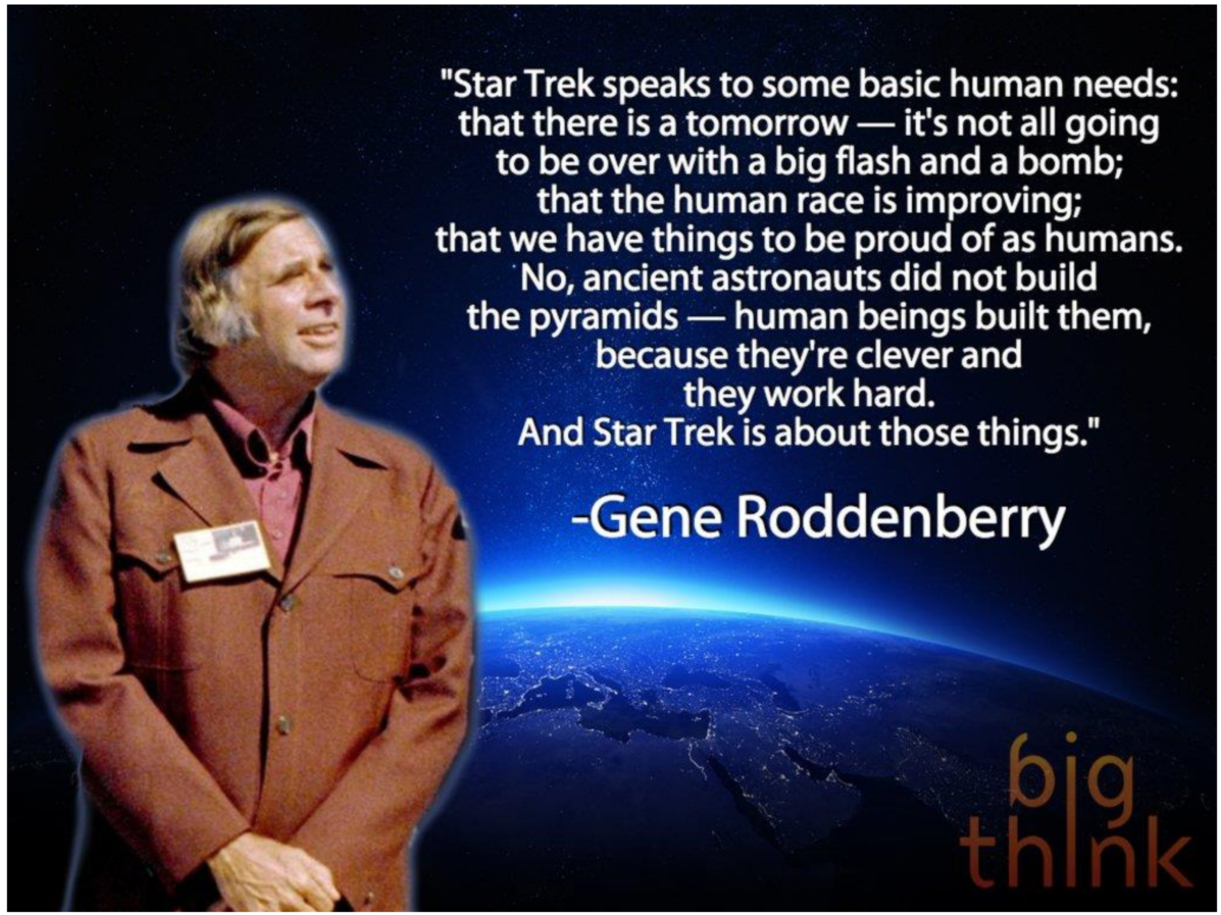

Die Vision von Kai-Fu Lee
Venture capitalist & computer scientist

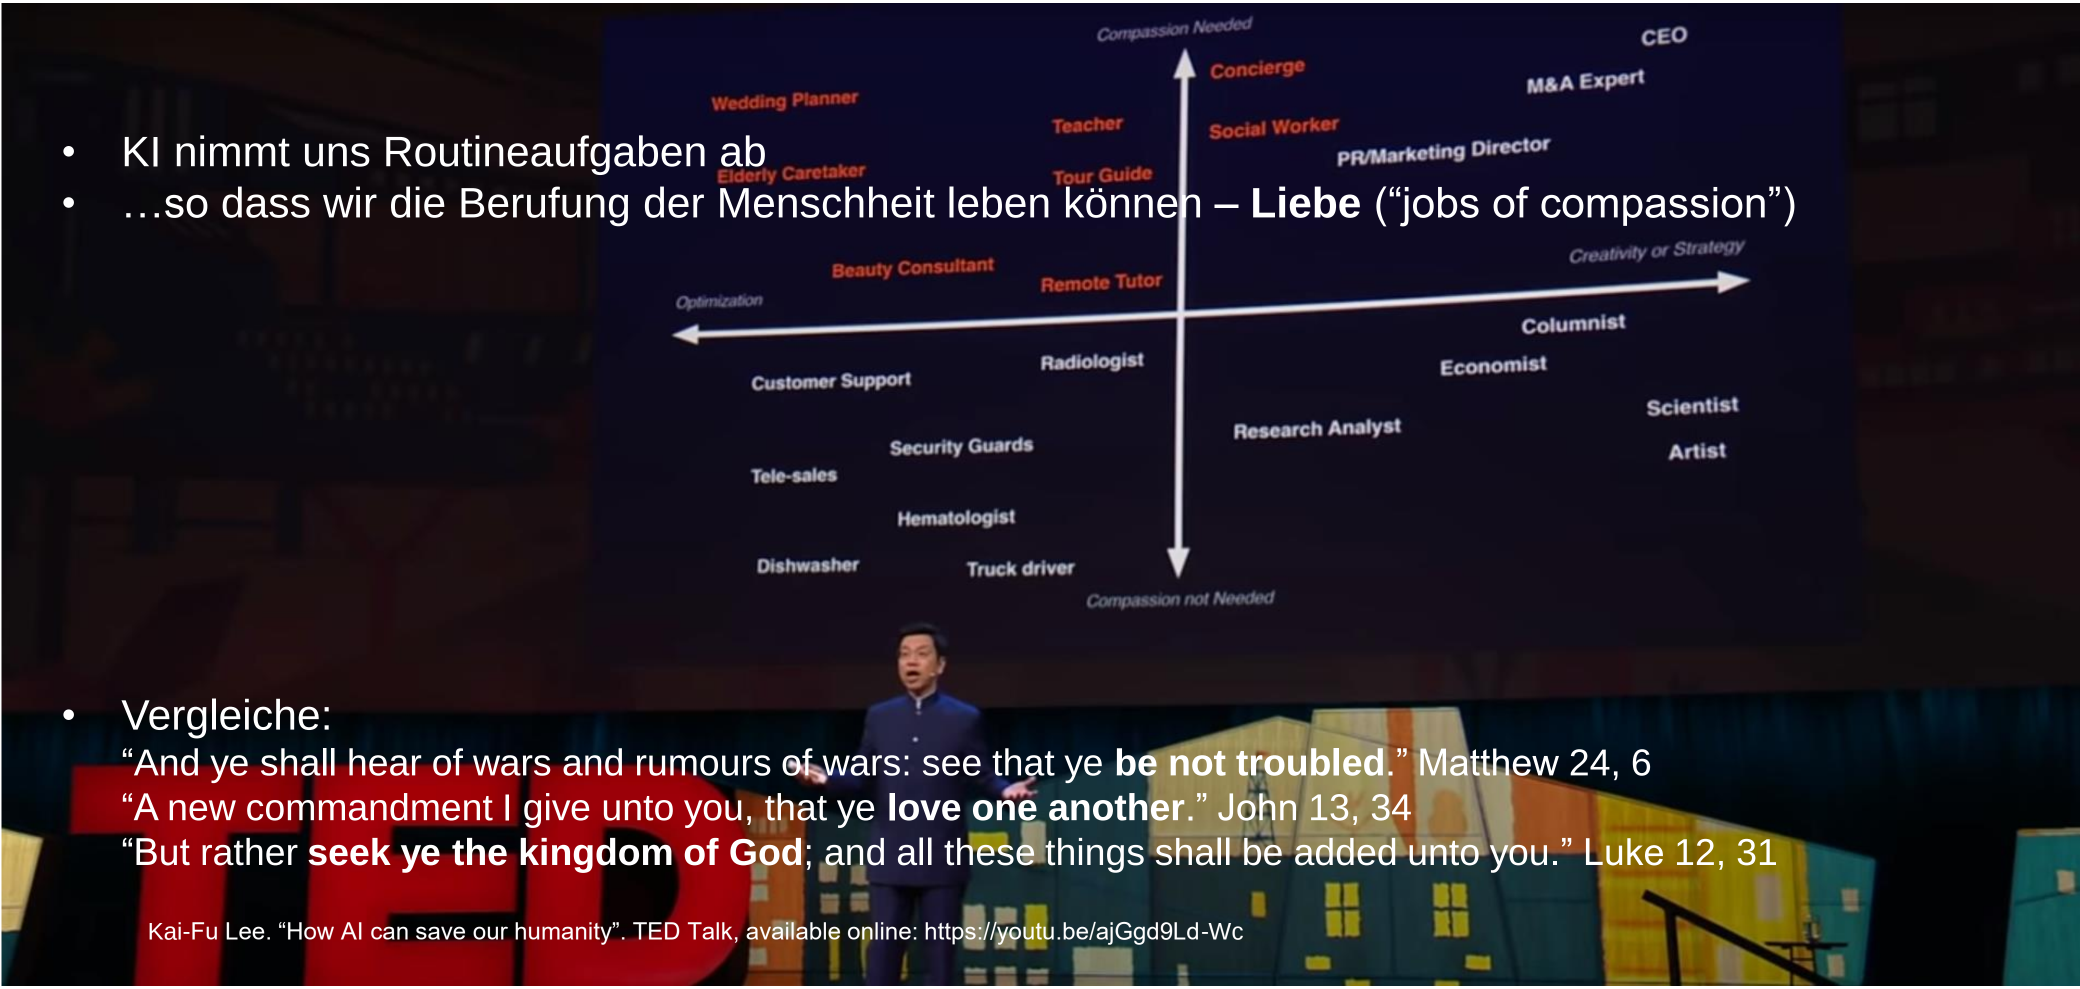

- KI nimmt uns Routineaufgaben ab
- …so dass wir die Berufung der Menschheit leben können – **Liebe** ("jobs of compassion")

- Vergleiche:
  "And ye shall hear of wars and rumours of wars: see that ye **be not troubled**." Matthew 24, 6
  "A new commandment I give unto you, that ye **love one another**." John 13, 34
  "But rather **seek ye the kingdom of God**; and all these things shall be added unto you." Luke 12, 31

Kai-Fu Lee. "How AI can save our humanity". TED Talk, available online: https://youtu.be/ajGgd9Ld-Wc

Zusammenfassung

Zürcher Hochschule
für Angewandte Wissenschaften

- Deep Learning wird *in normalen* Unternehmen **angewendet**
- Das *wirtschaftliche* Potential der Technologie wird zu **massenhafte**r **Verbreitung** führen
- Dies wird einen **grossen Wandel** in unsere Gesellschaften hervorrufen
- Hauptaufgabe: **guter Umgang** & Dialog miteinander (nicht technologischer Art – Liebe)

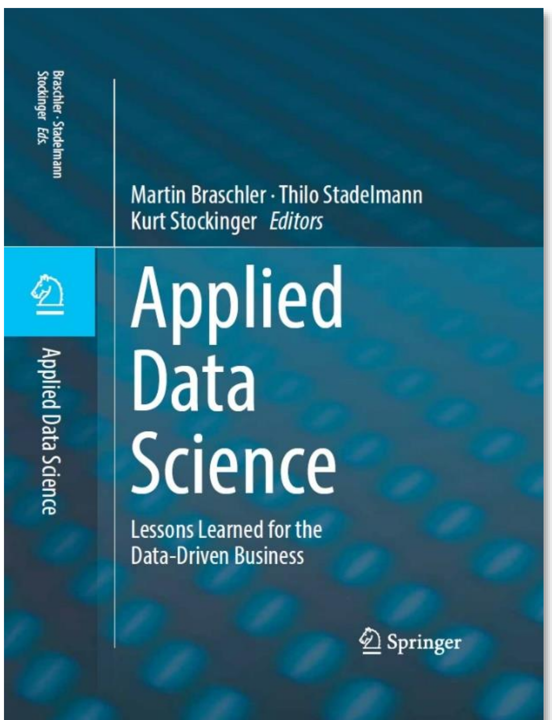

Über mich:

- Prof. AI/ML, scientific director ZHAW digital, head ZHAW Datalab, board Data+Service
- thilo.stadelmann@zhaw.ch
- +41 58 934 72 08
- @thilo_on_data
- https://stdm.github.io/

Weitere Kontakte:

- Data+Service Alliance: www.data-service-alliance.ch
- Zusammenarbeit: datalab@zhaw.ch

➔ Ich freue mich auf die Diksussion und generell Kontakt.

datalab
www.zhaw.ch/datalab

Swiss Alliance for
Data-Intensive Services

ANHANG

Pragmatisches Designparadigma: Rationale Agenten

Agents

- an **entity that perceives and acts**
- a **function from percept** histories **to actions** $f: P^* \to A$

Rational agents

- **For any** given class of **environn** **ek** the agent (or cl gents) with the **best performance**

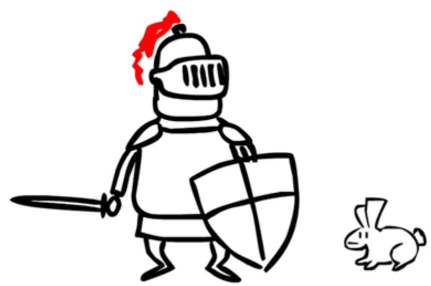

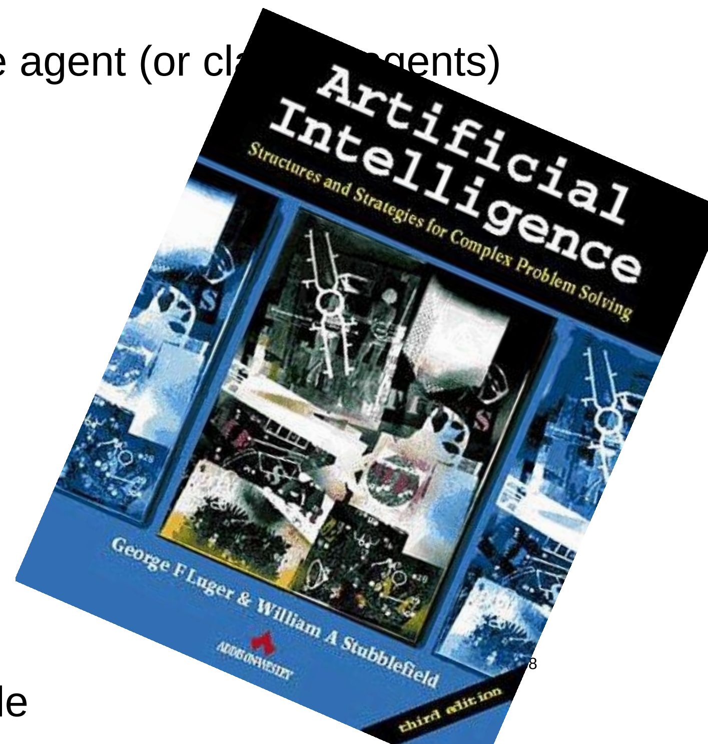

Caveat

- Computational limitations make perfect rationality unachievable

Ein Modell für generelle KI

Inspired by E. Mogenet @ Zurich ML Meetup #31

AI Knowledge engineering (symbolic):

- ↓ Ontologies
- ↓ Logical inference

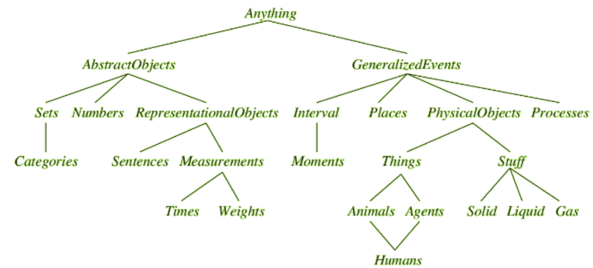

Gap to be filled by: **common sense DB**, **NLP**

Machine Learning (sub-symbolic):

- ↑ Hierarchical unsupervised learning
- ↑ Solid computer vision stack
- ↑ Images of the world

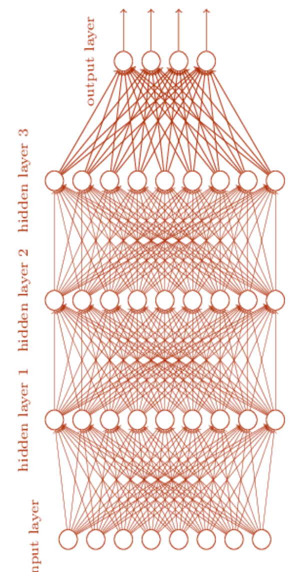

Was kann KI bereits heute?

1. Play a decent game of **table tennis** — ok
2. **Drive** safely along a curving **mountain road** — ok
3. Drive safely along **Technikumstrasse** Winterthur — ok (only since recently)
4. **Buy** a week's worth of **groceries on the web** — ok
5. Buy a week's worth of groceries **at Migros** — no
6. **Play** a decent game of **bridge** — ok
7. **Discover** and prove a new mathematical **theorem** — not complet
8. **Design** and execute a **research program** in molecular biology — not complet
9. Write an **intentionally funny** story — no
10. Give competent **legal advice** in a specialized area of law — ok
11. **Translate** spoken English **into spoken** Swedish in real time — ok
12. **Converse** successfully with another person for an hour — no
13. Perform a complex **surgical operation** — not complet
14. **Unload** any **dishwasher** and put everything away — no
15. Compete in the game show **Jeopardy!** — ok
16. **Write clickbait** articles fully automatized — ok
17. **Write math**ematical articles fully automatized — not completely

Gefahren durch KI?

- KI ist per Definition eine “**dual use Technology**”
  → siehe Report von Brundage et al., 2018
- Aber: “**natürliche Dummheit**” ist die grössere Bedrohung
- **Algorithmische Ethik** und **erklärbare KI** sind in den letzten Jahren zu einem top Forschungsfeld geworden – nicht wegen der unkalkulierbaren Risiken per se, sondern:

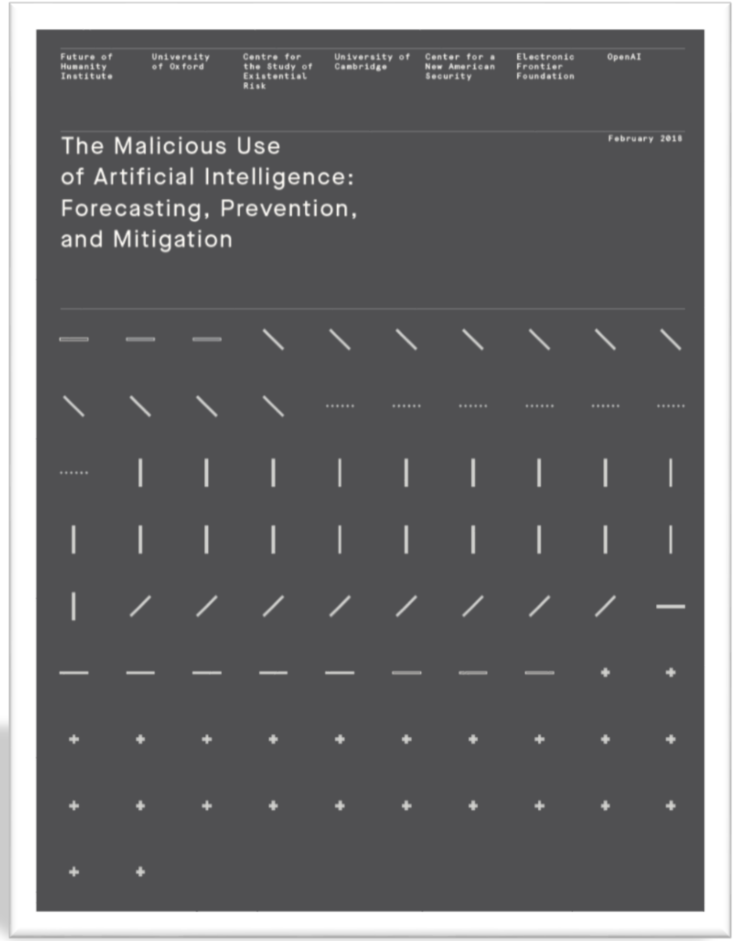

Developing for algorithmic fairness

The FAT ML code of conduct

See http://www.fatml.org/resources/principles-for-accountable-algorithms

FAT / ML

Purpose

- Help developers to **build algorithmic systems in publicly accountable ways**
- Accountability: the **obligation to report, explain, or justify** algorithmic decision-making & **mitigate** any **negative** social **impacts** or potential harms

Premise

- *A **human ultimately responsible** for decisions made/informed by an algorithm*

Principles

- **Responsibility**, **Explainability**, **Accuracy**, **Auditability**, **Fairness**

Make available somebody who will take care of adverse individual / societal effects	Explain any **algorithmic decision** in non-technical terms to end users	**Report** all **sources of uncertainty** / error in algorithms & data	Enable 3rd parties to **probe & understand** system **behavior**	Ensure algorithmic **decisions are not discriminatory** w.r.t. to people groups

Making it actionable

- **Publish** a Social Impact Statement
- …use above **principles as a guiding structure**
- …**revisit three times** during development process: design stage, pre-launch, post-launch

Suche der Parameter *einer Funktion*?

Neuron

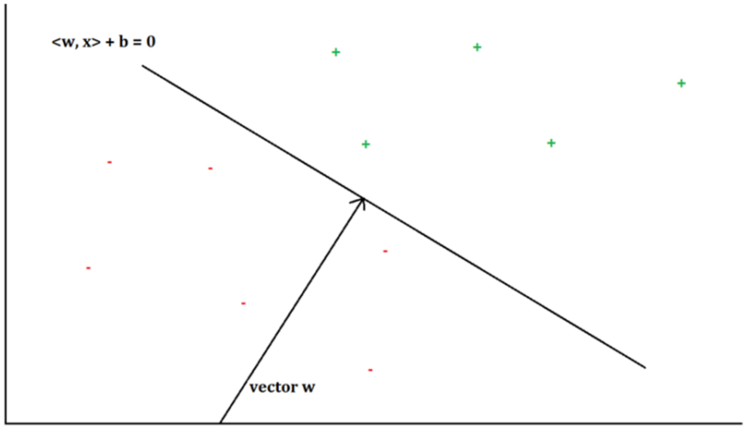

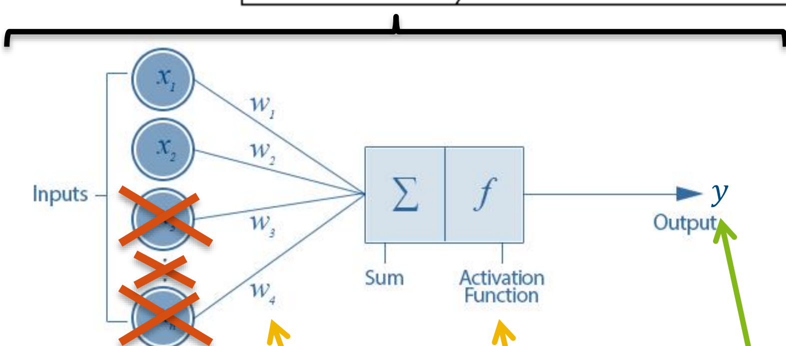

Merkmale (z.B. Pixel)

Anpassbare Parameter

Entscheidung (Schwellwert)

Ergebnis (z.B. «Auto»)

Neuronales Netz

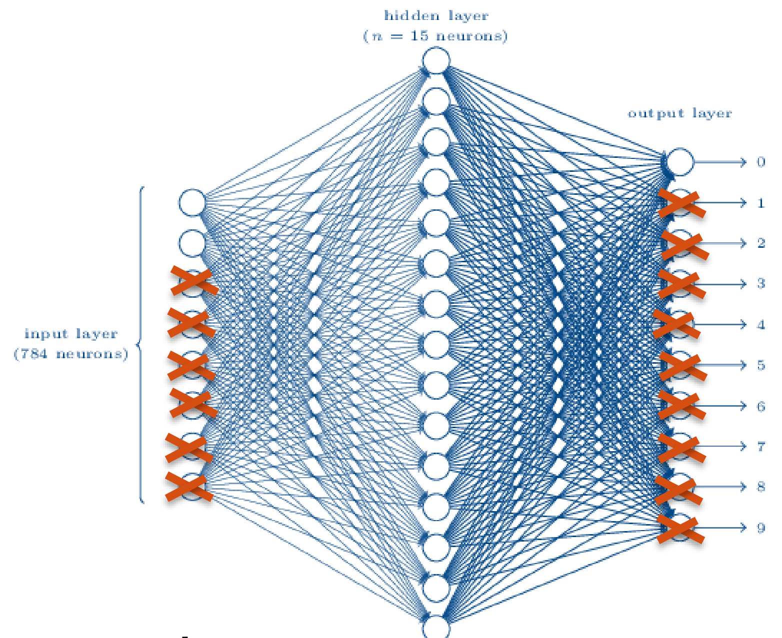

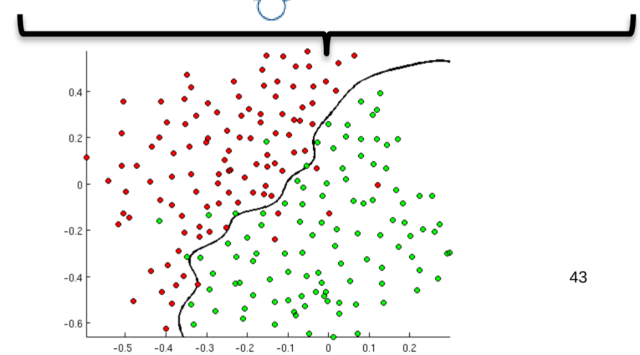

Suche der Parameter einer Funktion?

Wahrscheinlichkeit [%] für bestimmtes Ergebnis

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

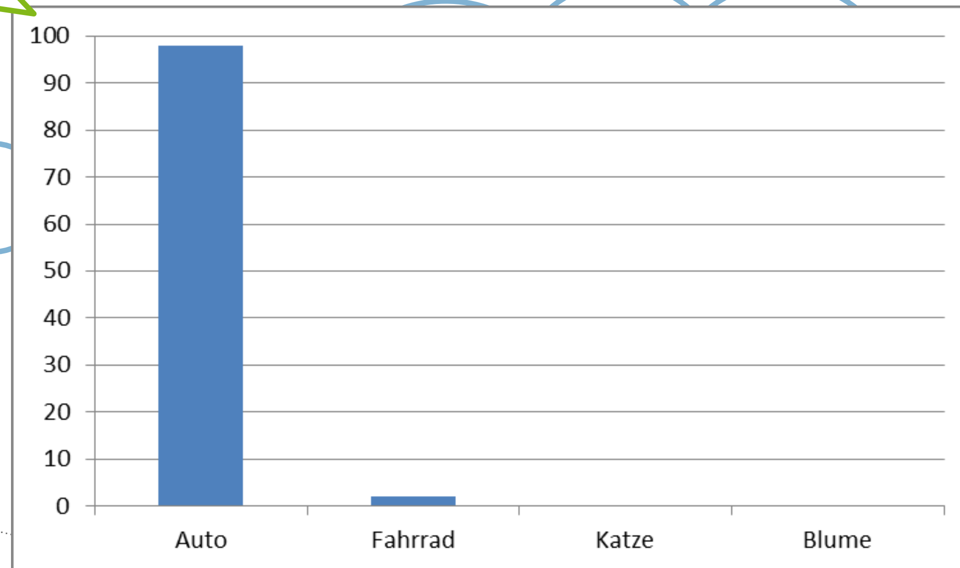

$$l(\boldsymbol{W}) = \underbrace{\frac{1}{N}\sum_{i=1}^{N}}_{\text{Durchschnitt (über alle Beispiele)}}\underbrace{(f_W(x_i) - y_i)}_{\text{Differenz IST – SOLL (Fehler)}}\underbrace{^2}_{\text{Bestraft grosse Fehler überproportional stärker}}$$

Durchschnitt (über alle Beispiele)

Differenz IST – SOLL (Fehler)

Bestraft grosse Fehler überproportional stärker

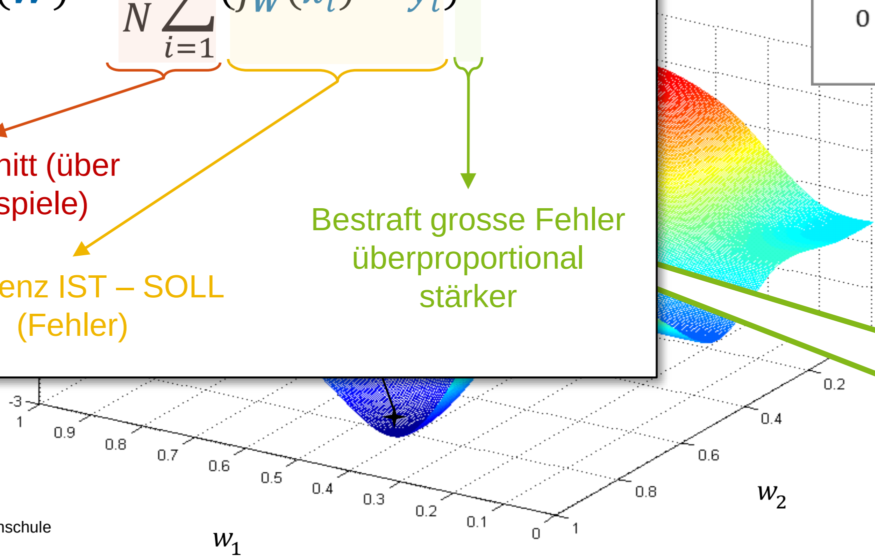

← Fehlerlandschaft

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

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

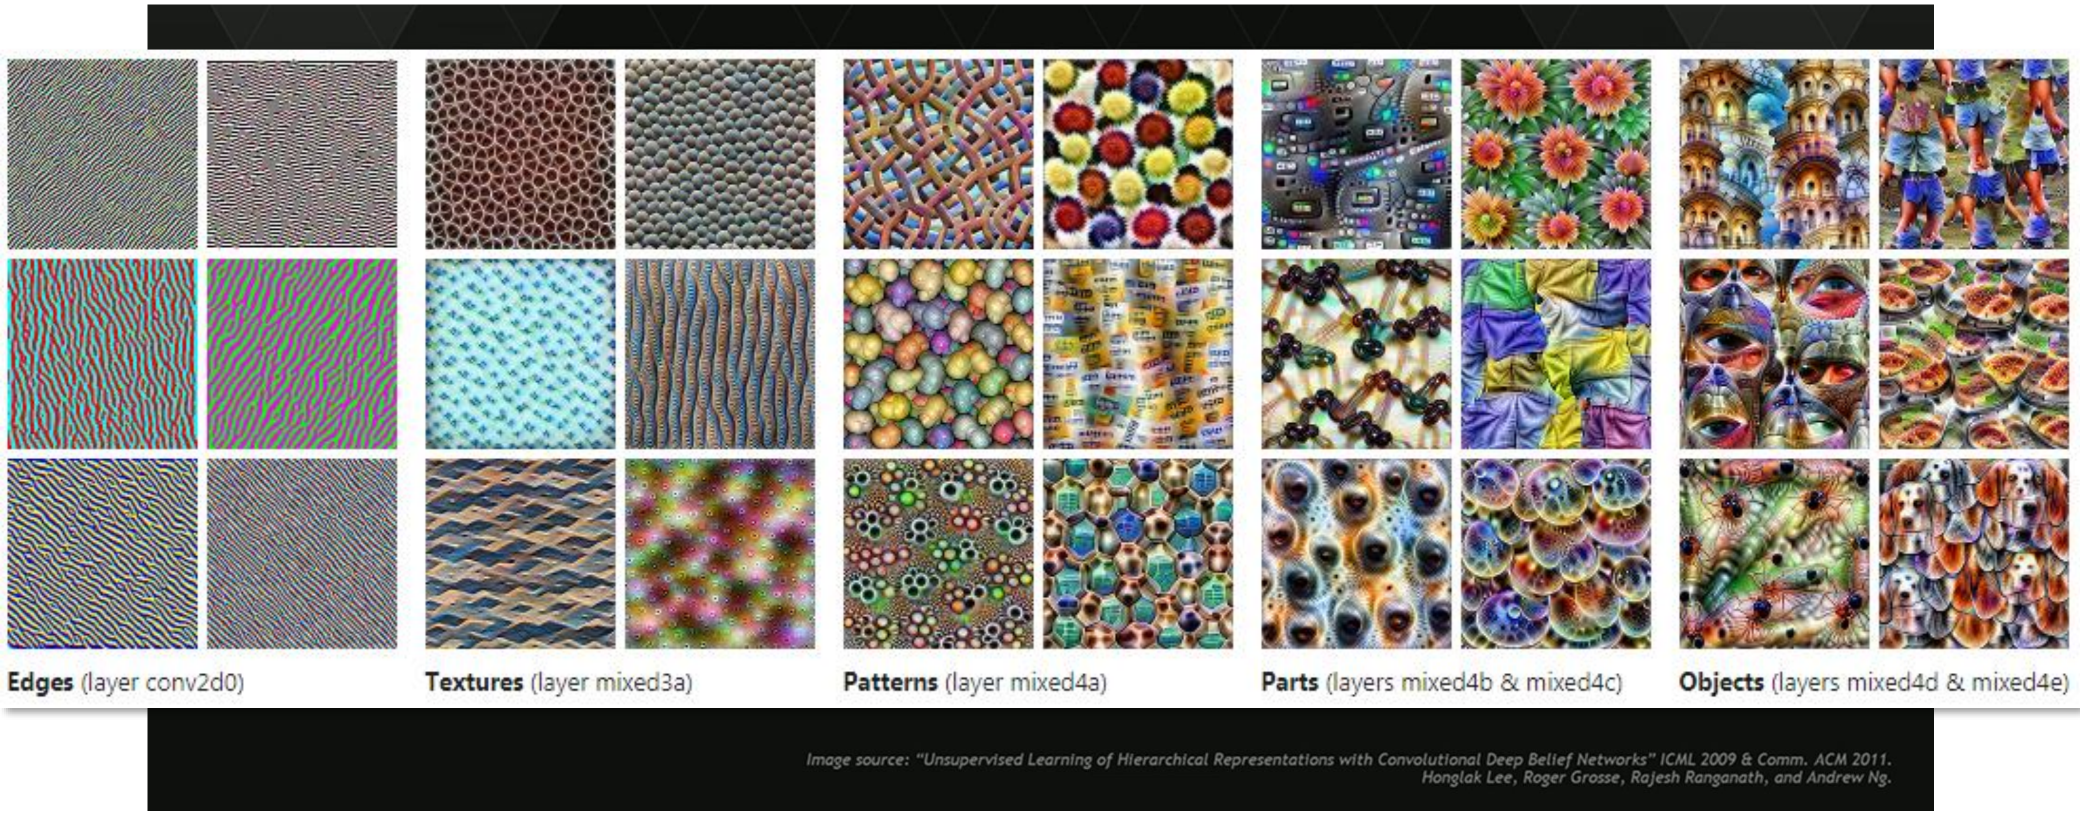

Quellen: https://www.pinterest.com/explore/artificial-neural-network/
Olah, et al., "Feature Visualization", Distill, 2017, https://distill.pub/2017/feature-visualization/.

Wie schlussfolgert die Maschine?

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

Verdeutlichen ein Problem:

- Adversarial Examples

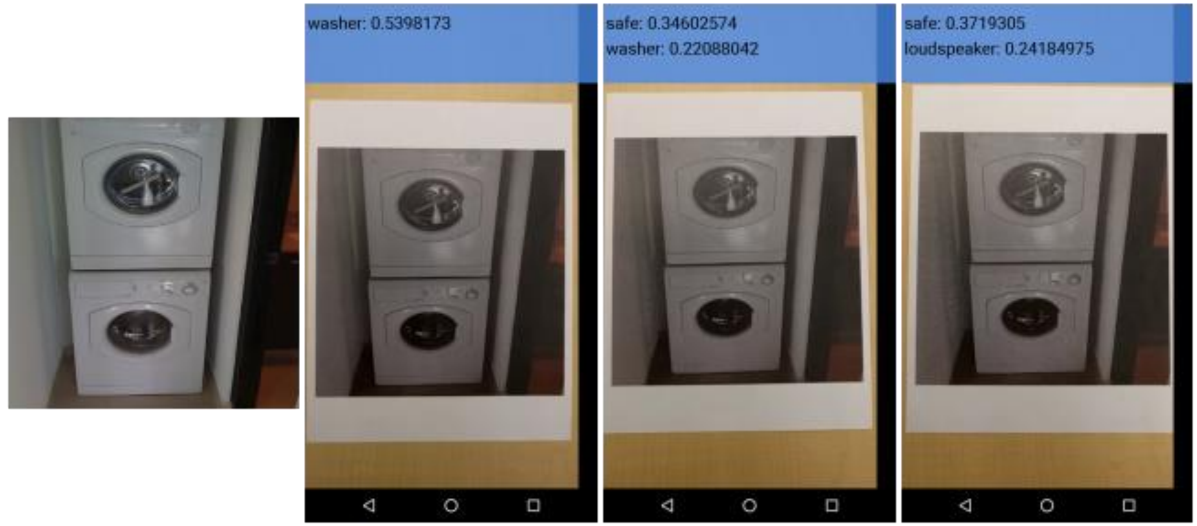

(a) Image from dataset (b) Clean image (c) Adv. image, $\epsilon = 4$ (d) Adv. image, $\epsilon = 8$

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

Bieten eine Lösung:

- Saliency Maps

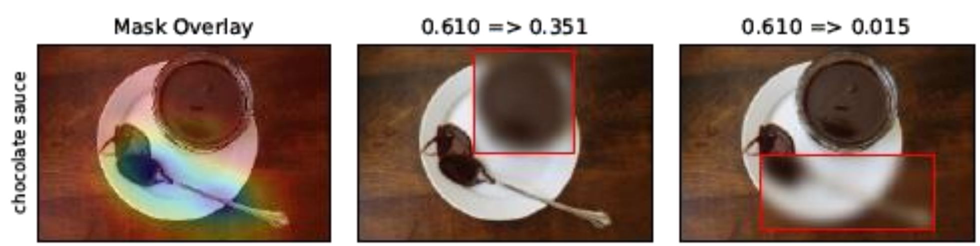

Ruth C. Fong & Andrea Vedaldi, «Interpretable Explanations of Black Boxes by Meaningful Perturbation», 2017

Trace & detect adversarial attacks
…using average local spatial entropy of feature response maps

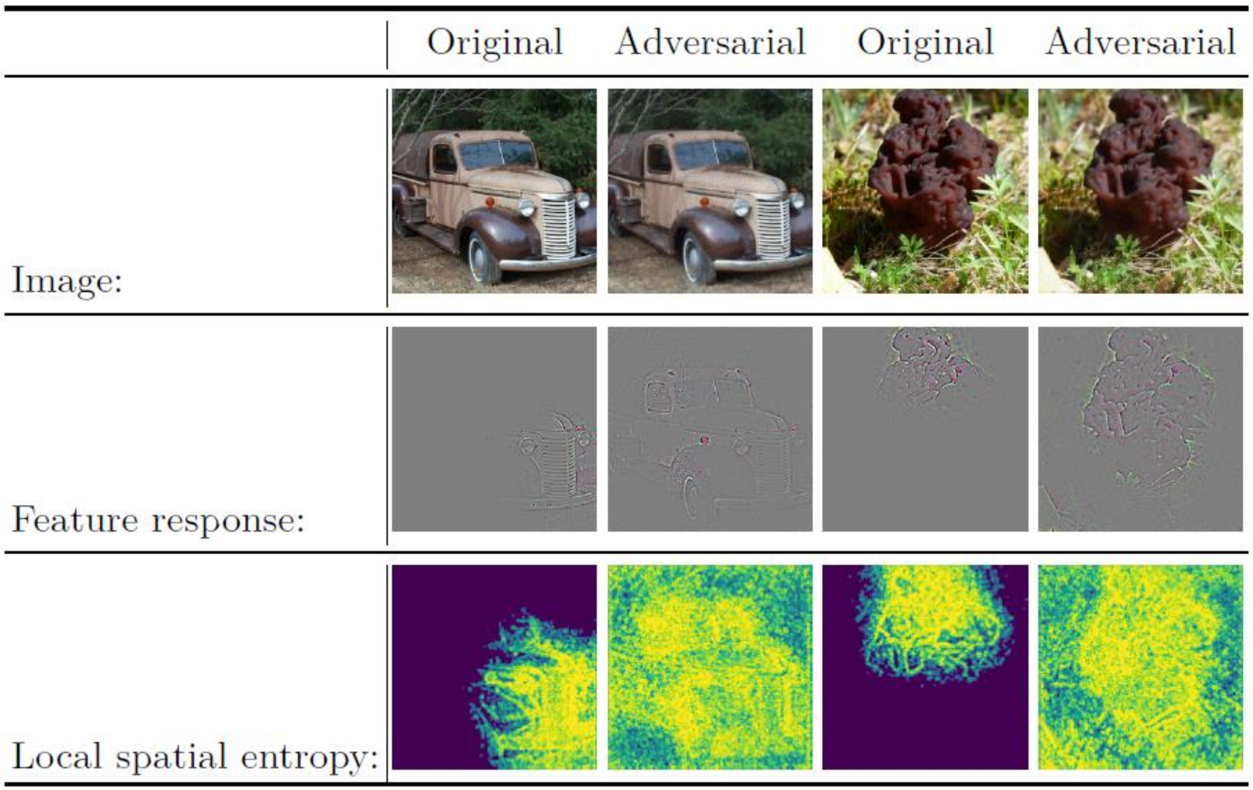

Amirian, Schwenker & Stadelmann (2018). *«Trace and Detect Adversarial Attacks on CNNs using Feature Response Maps»*. ANNPR'2018.

Zürcher Hochschule
für Angewandte Wissenschaften

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.

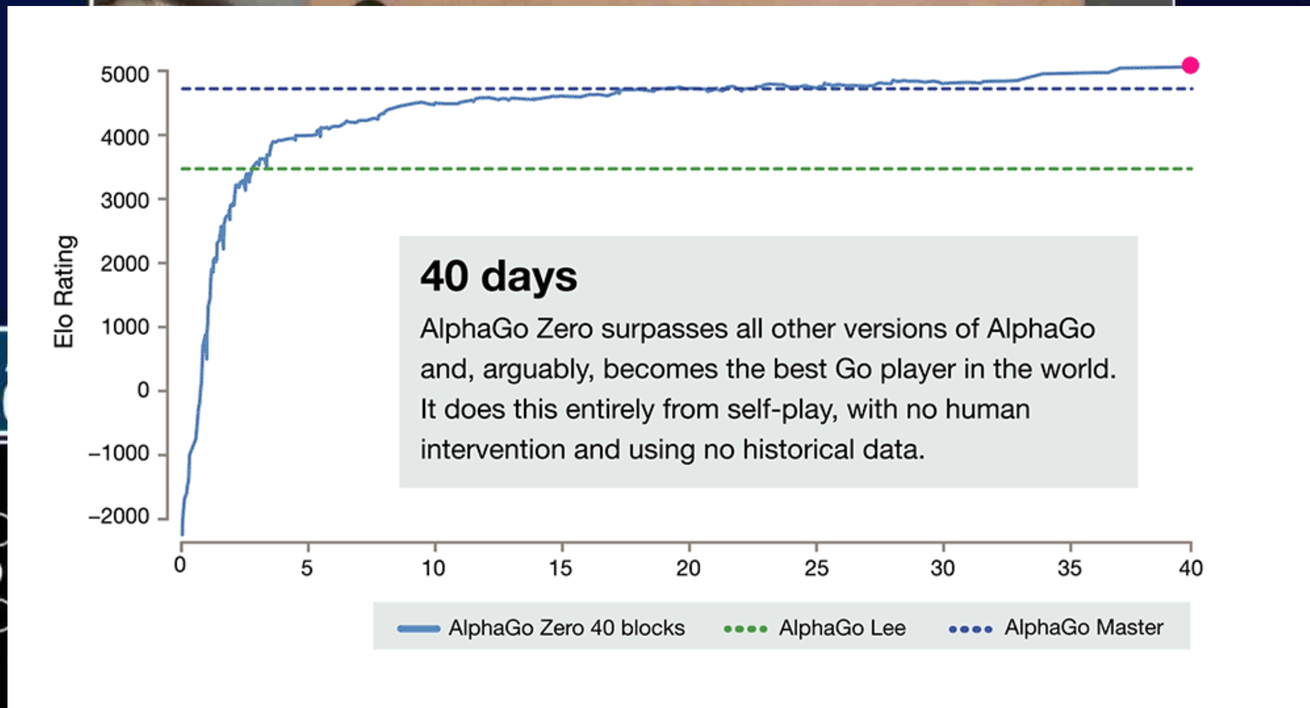

Deep neural networks can now transfer the style of one photo onto another

And the results are impressive

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

SHARE TWEET LINKEDIN

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

MIT Techn
Reviev

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

Künstliche Intelligenz

WaveNet lä[...]t Computersprache natürlich klingen

von Henning Steier / 12.9.201[...]

Die Google-Tochter DeepM[...]
macht auch Musik.

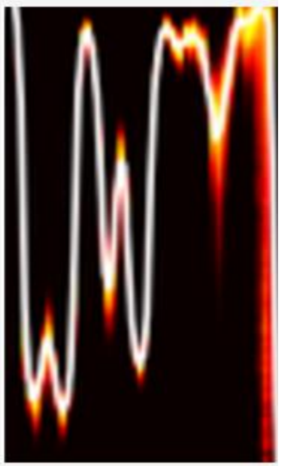

DeepMind lässt WaveNet Spra[...]

Die Google-Tochter Deep[...]
Spiel «Go» Schlagzeilen: [...]
einen der besten mensch[...]
Londoner Unternehmen [...]
erzeugt Sprache, die sehr [...]
im Blogeintrag des Unter[...]
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.

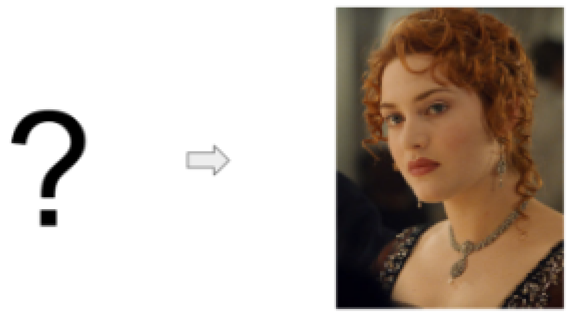

Model Architecture

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.

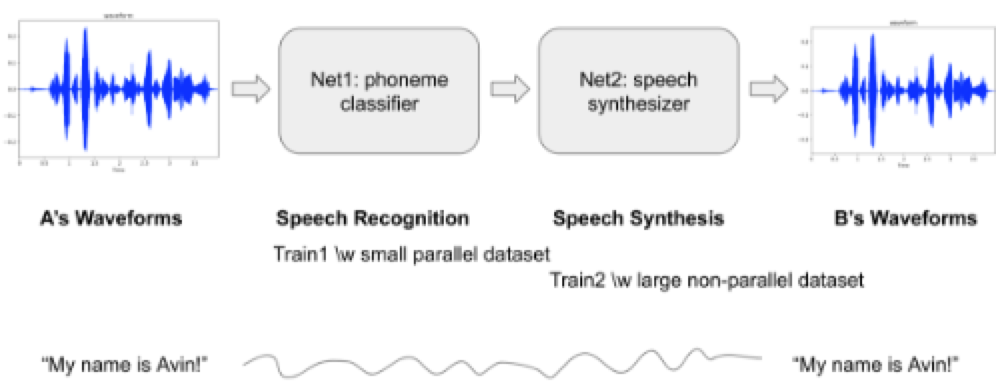

Zürcher Hochschule
für Angewandte Wissenschaften

[...]nerierte Sprache
[...]s Texteingabe»

[...]nerierte Musik
[...]ne Inhaltsvorgabe»

Zürcher Hochschule
für Angewandte Wissenschaften

1. ML @ Information Engineering Group

Institute of Applied Information Technology, School of Engineering

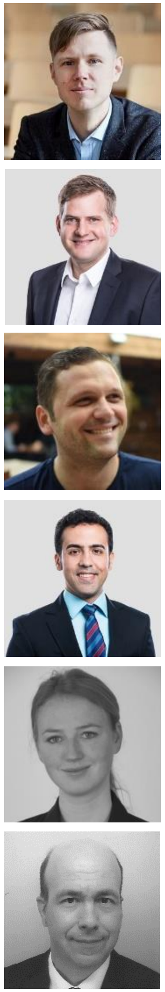

Machine learning-based Pattern Recognition

- Robust Deep Learning
- Voice Recognition
- Document Analysis
- Learning to Learn & Control

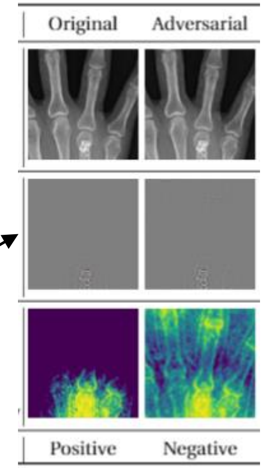

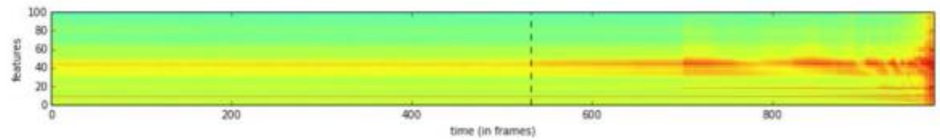

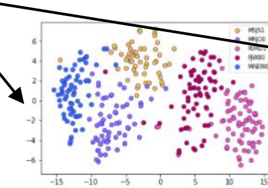

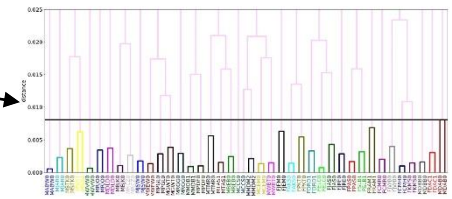

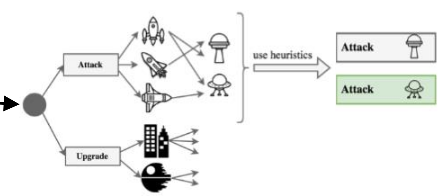

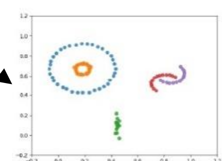

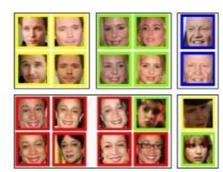

1. Face matching

Zürcher Hochschule
für Angewandte Wissenschaften

1. Face matching – challenges & solutions

[!] DEEPIMPACT

Schweizerische Eidgenossenschaft
Confédération suisse
Confederazione Svizzera
Confederaziun svizra

Swiss Confederation

Innosuisse – Swiss Innovation Agency

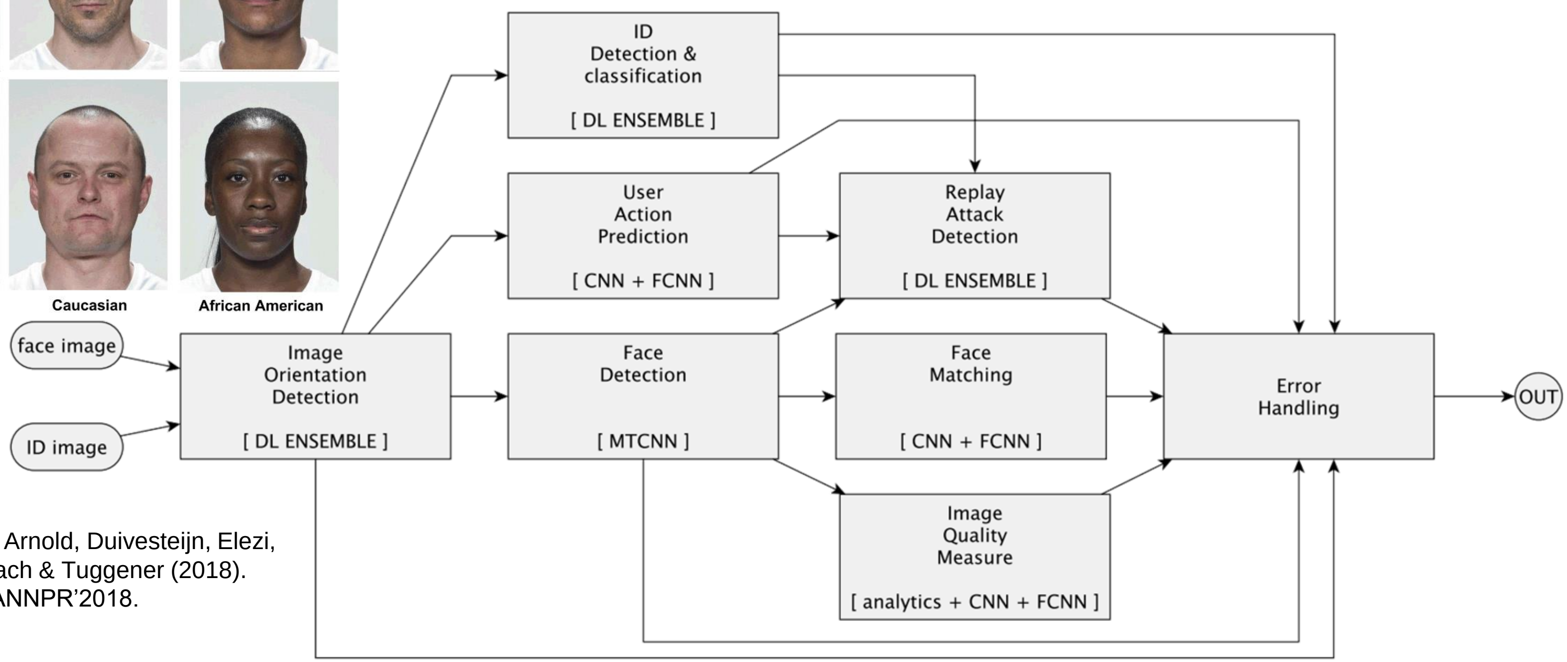

Stadelmann, Amirian, Arabaci, Arnold, Duivesteijn, Elezi, Geiger, Lörwald, Meier, Rombach & Tuggener (2018). *«Deep Learning in the Wild»*. ANNPR'2018.

Zürcher Hochschule
für Angewandte Wissenschaften

2. Print media monitoring

Task


International.
Basler Zeitung | Mittwoch, 20. Juni 2018 | Seite 6
Nachrichten
Spionage für den Erzfeind Iran
Israelischer Ex-Minister arbeitete als Agent für die Mullahs. Jetzt droht ihm lebenslänglich
Vermögen beschlagnahmt
Asylbewerber können bleiben
Europäischer Gerichtshof relativiert Ablehnungsbescheide
Klage von Le Pen abgelehnt
Nordkoreanischer Diktator zu Besuch in Peking

Challenge


Mittwoch, 20. Juni 2018
Sport Blick 15
Sein Juniorentrainer Mano Pavesi über unseren WM-Helden Steven Zuber
«Steven hat sich alles selber beigebracht»
Unsere Nati ist gegen Serbien unter Druck
Verlieren verboten!
Transfer Ticker
Liverpool will Yann Sommer!

Nuisance


Freitag, 22. Juni 2018
Das Tages-Horoskop
Liebling der Sterne
Löwe
Blick 23
15,1 Millionen
Sind Sie der nächste Lotto-König?
Jungfrau
Waage
Skorpion
Schütze
Steinbock
Wassermann
Fische
Widder
Stier
Zwillinge
Krebs
Schweden-Rätsel
Gesamtwert: 5000 Franken!
Wochenpreis: 1 x sieben Nächte für 2 Personen, inkl. HP, im ****Seehotel Pilatus Hergiswil im Wert von 3000 Franken!
Sudoku
Flottenmanöver

ARGUS DATA INSIGHTS
WISSEN ZUM ERFOLG

Schweizerische Eidgenossenschaft
Confédération suisse
Confederazione Svizzera
Confederaziun svizra

Swiss Confederation

Innosuisse – Swiss Innovation Agency

Zürcher Hochschule
für Angewandte Wissenschaften

2. Print media monitoring – ML solution

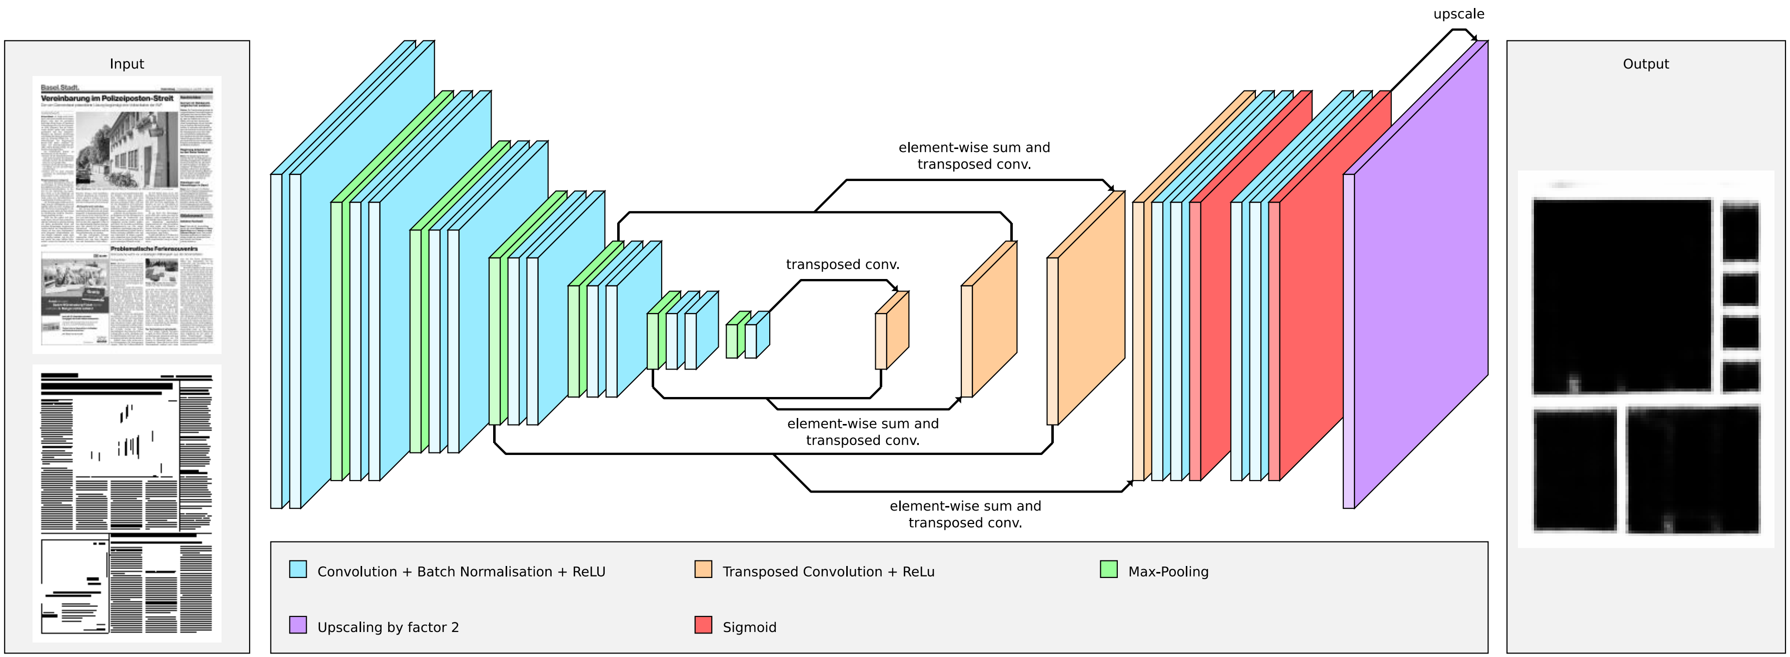

Meier, Stadelmann, Stampfli, Arnold & Cieliebak (2017). *«Fully Convolutional Neural Networks for Newspaper Article Segmentation»*. ICDAR'2017.
Stadelmann, Tolkachev, Sick, Stampfli & Dürr (2018). *«Beyond ImageNet - Deep Learning in Industrial Practice»*. In: Braschler et al., «Applied Data Science», Springer.

3. Industrial quality control

Schweizerische Eidgenossenschaft
Confédération suisse
Confederazione Svizzera
Confederaziun svizra

Swiss Confederation

Innosuisse – Swiss Innovation Agency

Task

- Reliably **sort out faulty balloon catheters** in image-based production quality control

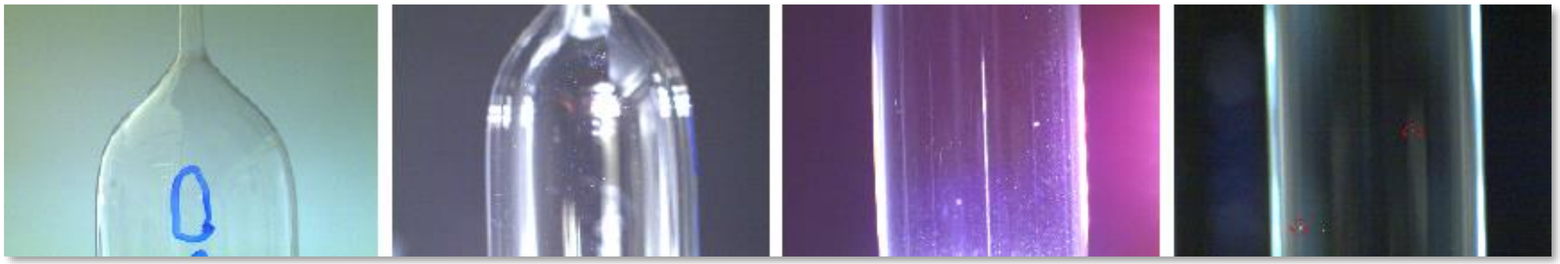

Challenges

- **Non-natural** image source, class **imbalance**, **optic**al conditions, **variation** in defect size & shape

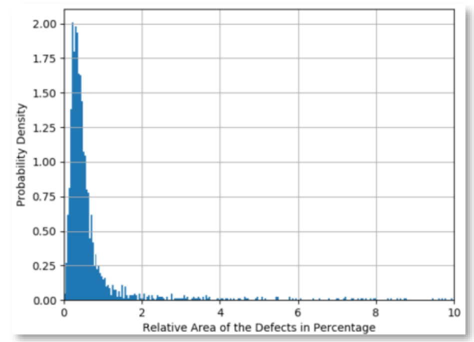

3. Industrial quality control – baseline results

Ingredients
- Weighted loss
- Defect cropping
- Careful customization

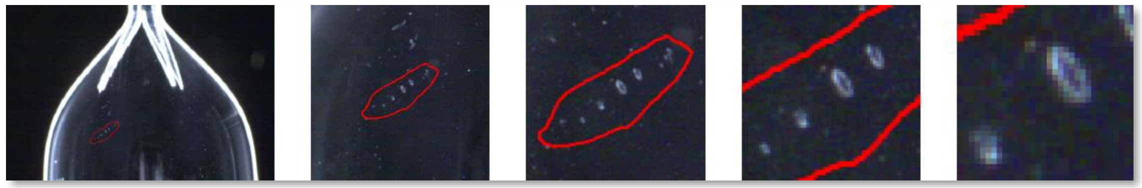

Interim results

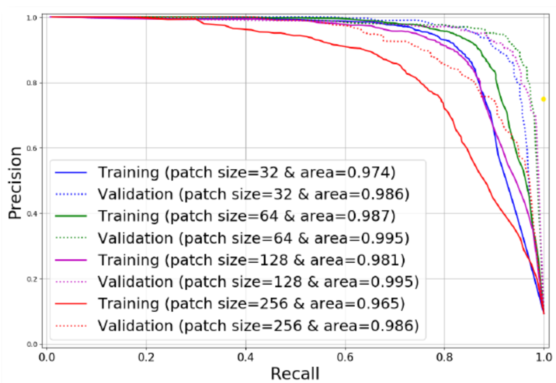

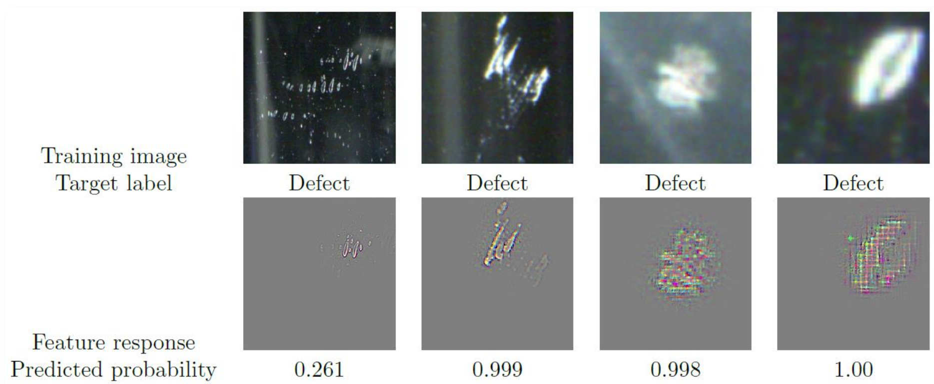

Zürcher Hochschule
für Angewandte Wissenschaften

3. Industrial quality control – recent results
(Work in progress)

- Human performance isn't flawless
- Tailoring pays off
- Data shortage may be outsmarted

BW-TEC

Schweizerische Eidgenossenschaft
Confédération suisse
Confederazione Svizzera
Confederaziun svizra

Swiss Confederation

Innosuisse – Swiss Innovation Agency

Defect

Good

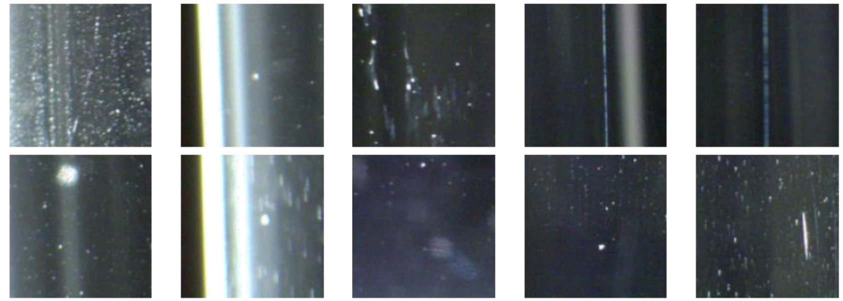

Figure 2: Samples of failure cases in classification. The shown *defect* samples in the table are not recognized as a defects, and the *good* images are misclassified as defects.

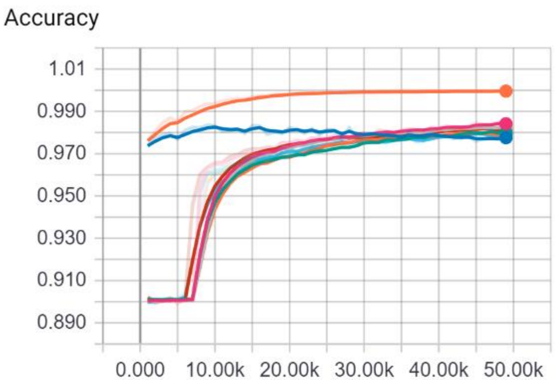

Name	Smoothed	Value	Step	Time	Relative
QualitAI_VGG19_Full_Pretrained\train	0.9996	0.9996	49.00k	Tue Jan 22, 02:32:13	8h 30m 56s
QualitAI_VGG19_Full_Pretrained\validation	0.9776	0.9783	49.00k	Tue Jan 22, 02:32:24	8h 30m 56s
QualitAI_VGG19_Full_Random\train	0.9841	0.9841	49.00k	Thu Jan 24, 19:28:02	10h 29m 2s
QualitAI_VGG19_Full_Random\validation	0.9798	0.9798	49.00k	Thu Jan 24, 19:28:14	10h 29m 2s
QualitAI_VGG19_Half\train	0.9827	0.9835	49.00k	Thu Jan 24, 13:01:47	4h 9m 12s
QualitAI_VGG19_Half\validation	0.9792	0.9798	49.00k	Thu Jan 24, 13:01:54	4h 9m 11s
QualitAI_VGG19_Quarter\train	0.9817	0.9823	49.00k	Thu Jan 24, 10:53:52	2h 17m 21s
QualitAI_VGG19_Quarter\validation	0.9791	0.9806	49.00k	Thu Jan 24, 10:53:56	2h 17m 21s

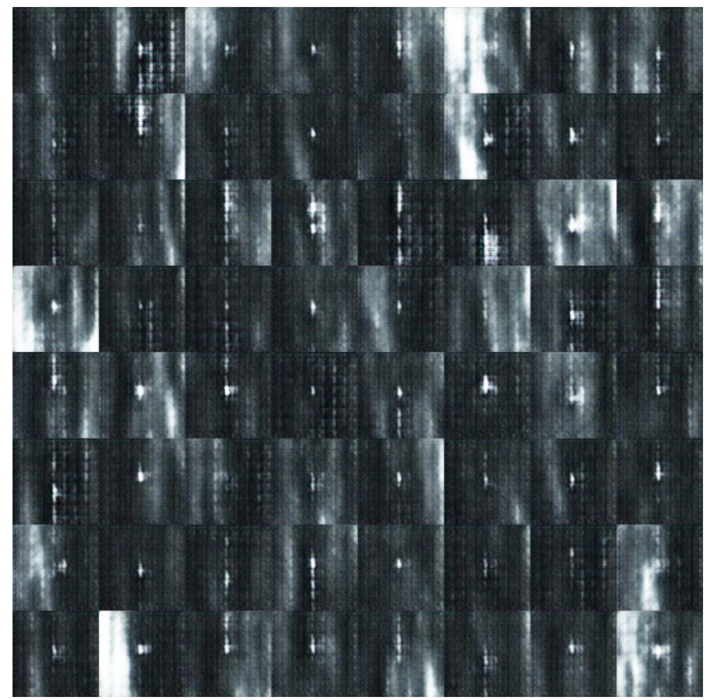

Zürcher Hochschule
für Angewandte Wissenschaften

SCOREPAD

Schweizerische Eidgenossenschaft
Confédération suisse
Confederazione Svizzera
Confederaziun svizra

Swiss Confederation

Innosuisse – Swiss Innovation Agency

4. Music scanning

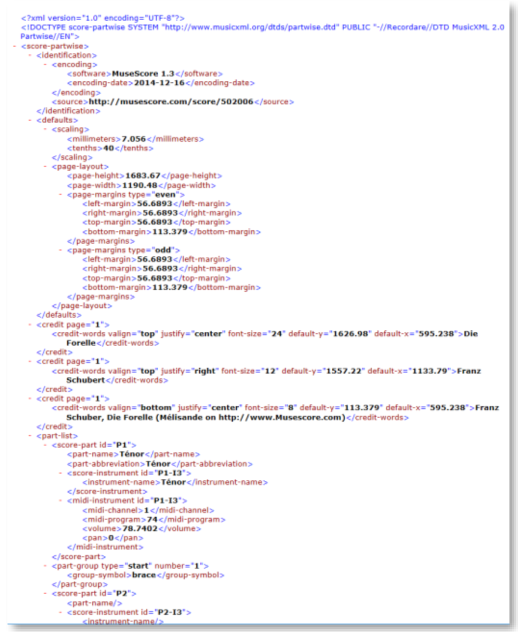

```
<?xml version="1.0" encoding="UTF-8"?>
<!DOCTYPE score-partwise SYSTEM "http://www.musicxml.org/dtds/partwise.dtd" PUBLIC "-//Recordare//DTD MusicXML 2.0 Partwise//EN">
- <score-partwise>
  - <identification>
    - <encoding>
        <software>MuseScore 1.3</software>
        <encoding-date>2014-12-16</encoding-date>
      </encoding>
      <source>http://musescore.com/score/502006</source>
    </identification>
  - <defaults>
    - <scaling>
        <millimeters>7.056</millimeters>
        <tenths>40</tenths>
      </scaling>
    - <page-layout>
        <page-height>1683.67</page-height>
        <page-width>1190.48</page-width>
      - <page-margins type="even">
          <left-margin>56.6893</left-margin>
          <right-margin>56.6893</right-margin>
          <top-margin>56.6893</top-margin>
          <bottom-margin>113.379</bottom-margin>
        </page-margins>
      - <page-margins type="odd">
          <left-margin>56.6893</left-margin>
          <right-margin>56.6893</right-margin>
          <top-margin>56.6893</top-margin>
          <bottom-margin>113.379</bottom-margin>
        </page-margins>
      </page-layout>
    </defaults>
  - <credit page="1">
      <credit-words valign="top" justify="center" font-size="24" default-y="1626.98" default-x="595.238">Die
        Forelle</credit-words>
    </credit>
  - <credit page="1">
      <credit-words valign="top" justify="right" font-size="12" default-y="1557.22" default-x="1133.79">Franz
        Schubert</credit-words>
    </credit>
  - <credit page="1">
      <credit-words valign="bottom" justify="center" font-size="8" default-y="113.379" default-x="595.238">Franz
        Schuber, Die Forelle (Mélisande on http://www.Musescore.com)</credit-words>
    </credit>
  - <part-list>
    - <score-part id="P1">
        <part-name>Ténor</part-name>
        <part-abbreviation>Ténor</part-abbreviation>
      - <score-instrument id="P1-I3">
          <instrument-name>Ténor</instrument-name>
        </score-instrument>
      - <midi-instrument id="P1-I3">
          <midi-channel>1</midi-channel>
          <midi-program>74</midi-program>
          <volume>78.7402</volume>
          <pan>0</pan>
        </midi-instrument>
      </score-part>
    - <part-group type="start" number="1">
        <group-symbol>brace</group-symbol>
      </part-group>
    - <score-part id="P2">
        <part-name/>
      - <score-instrument id="P2-I3">
          <instrument-name/>
```


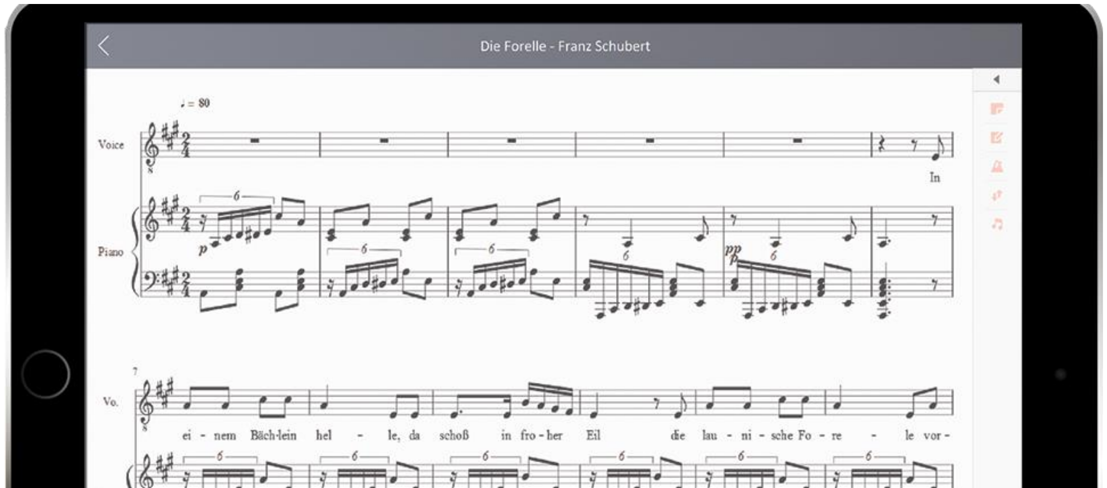

4. Music scanning – challenges & solutions

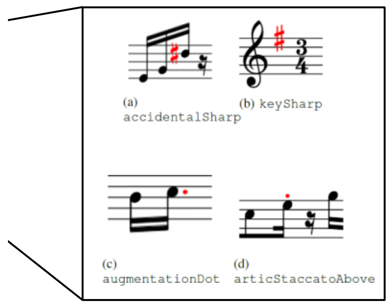

,

SCOREPAD

Schweizerische Eidgenossenschaft
Confédération suisse
Confederazione Svizzera
Confederaziun svizra

Swiss Confederation

Innosuisse – Swiss Innovation Agency

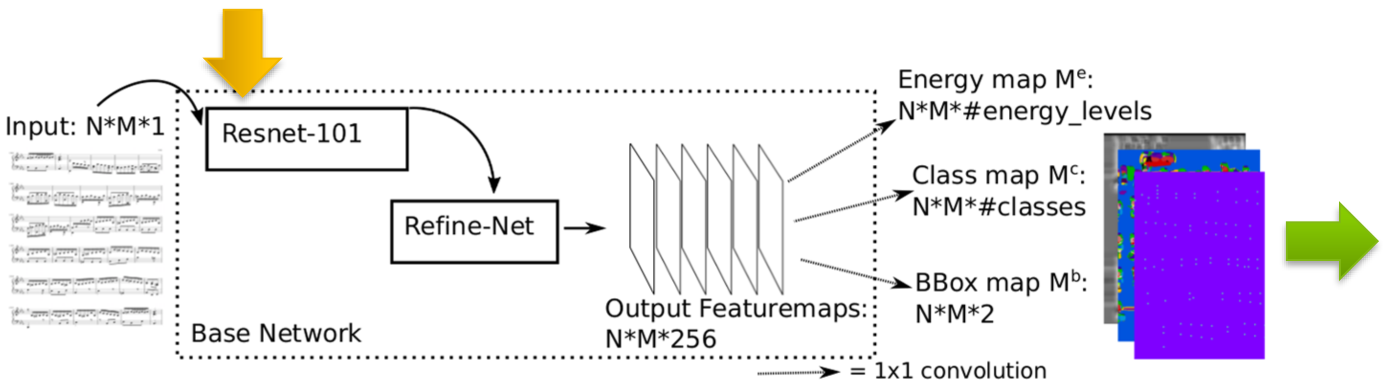

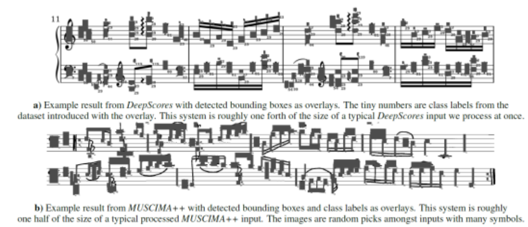

Tuggener, Elezi, Schmidhuber, Pelillo & Stadelmann (2018). *«DeepScores – A Dataset for Segmentation, Detection and Classification of Tiny Objects»*. ICPR'2018.
Tuggener, Elezi, Schmidhuber & Stadelmann (2018). *«Deep Watershed Detector for Music Object Recognition»*. ISMIR'2018.

Zürcher Hochschule für Angewandte Wissenschaften

4. Music scanning – industrialization

SCOREPAD

Schweizerische Eidgenossenschaft
Confédération suisse
Confederazione Svizzera
Confederaziun svizra

Swiss Confederation

Innosuisse – Swiss Innovation Agency

Recent results on class imbalance and robustness challenges

1. Added sophisticated **data augmentation** in every page's margins

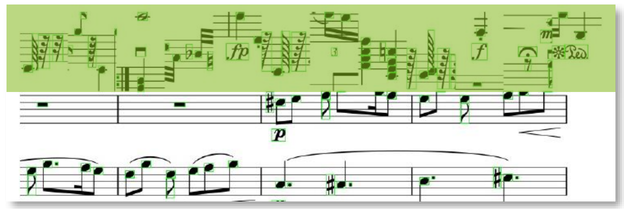

2. Put additional effort (and compute) into hyperparameter **tuning** and **longer training**
3. Trained also on scanned (more **real-worldish**) scores

➔ **Improved** our **mAP** from 16% (on purely synthetic data) **to 73%** on more challenging real-world data set (additionally, using Pacha et al.'s evaluation method as a 2nd benchmark: from 24.8% to 47.5%)

Elezi, Tuggener, Pelillo & Stadelmann (2018). *«DeepScores and Deep Watershed Detection: current state and open issues»*. WoRMS @ ISMIR'2018.
Pacha, Hajic, Calvo-Zaragoza (2018). *«A Baseline for General Music Object Detection with Deep Learning»*. Appl. Sci. 2018, 8, 1488, MDPI.

The vision of Jesus Christ

*"And ye shall hear of wars and rumours of wars: **see that ye be not troubled**."*
Matthew 24, 6

*"A new commandment I give unto you, that
ye **love one another**."*
John 13, 34

*"But **rather seek ye** the kingdom of God **[things above]**;
and all these things shall be added unto you."*
Luke 12, 31 [Colossians 3, 2]