Zürcher Hochschule
für Angewandte Wissenschaften

Wie künstliche Intelligenz und maschinelles Lernen unser Leben verändern

Marketing Arena Schaffhausen 2019: «Digitale Transformation – The Big Five», 17. September 2019

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 scientific field concerned with complex computer applications

"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 gehört zu künstlicher Intelligenz?

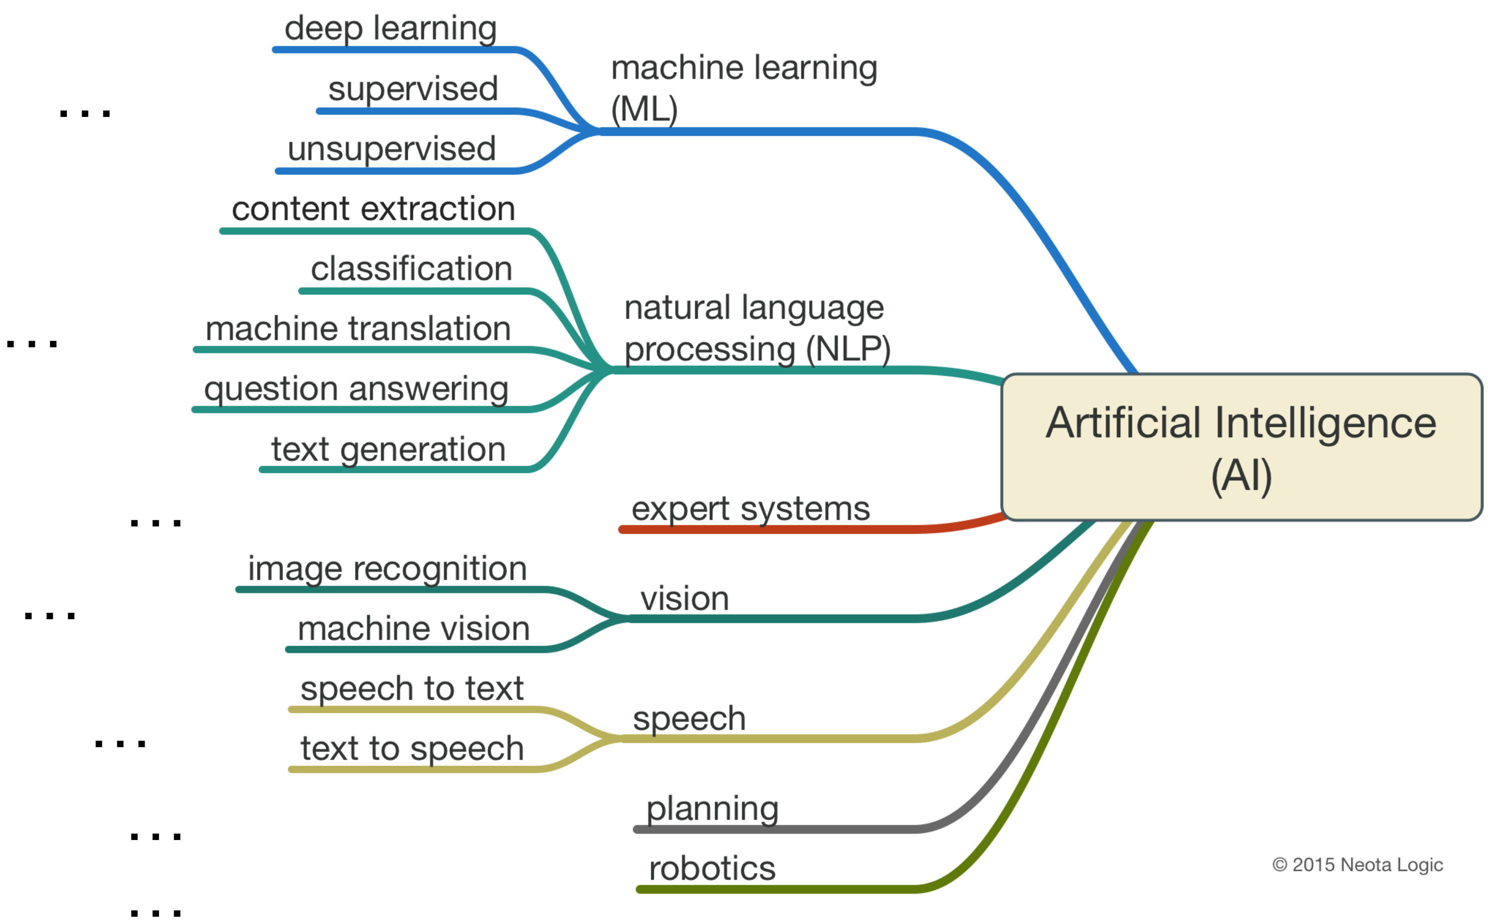

Eine kurze Geschichte der KI
Wiederkehrende Muster überzogener Erwartungen

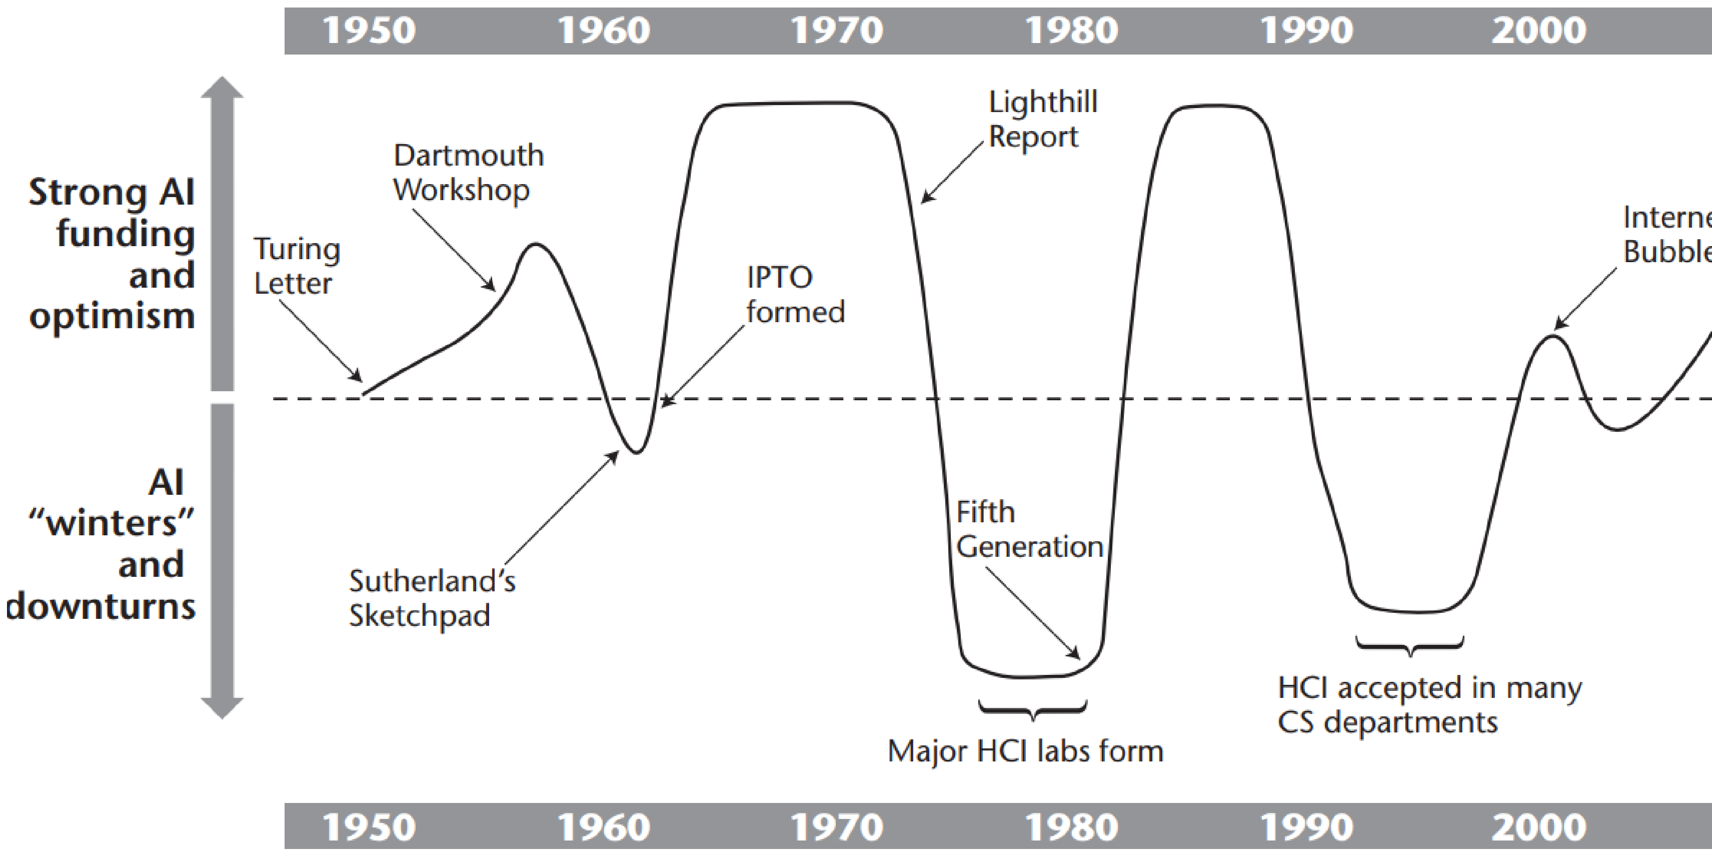

Quelle: Grudin, "AI and HCI: Two Fields Divided by a Common Focus", AI Magazine, Winter 2009.

Was ist passiert?
Der ImageNet Wettbewerb

1000 Kategorien
1 Mio. Beispiele

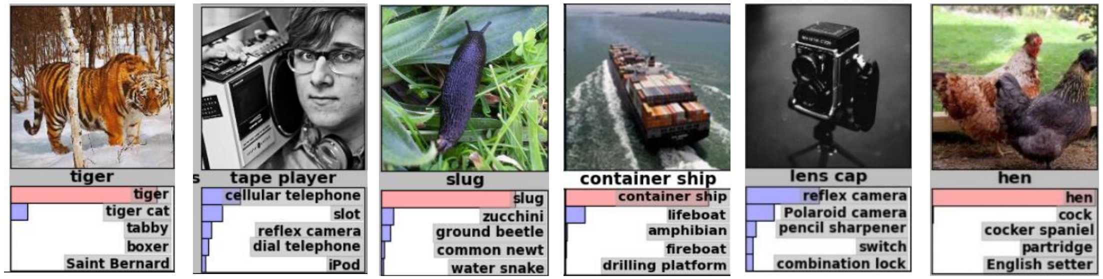

…

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.)

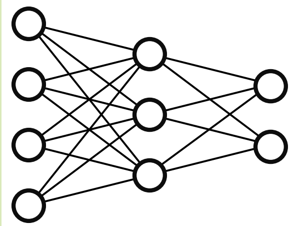

Kontainerschiff

Tiger

Grundidee Deep Learning: “feature learning”

Merkmalsextraktion manuell definierter Deskriptoren

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

Bildklassifikation (herkömmlicher Ansatz)

(0.2, 0.4, …)

(0.4, 0.3, …)

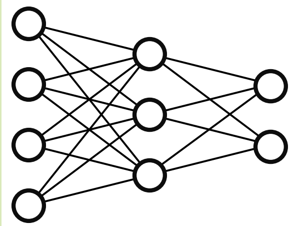

Kontainerschiff

Tiger

Rohdaten als Input, wesentliche Merkmale werden automatisch gelernt

Bildklassifikation (neu: Convolutional Neural Networks)

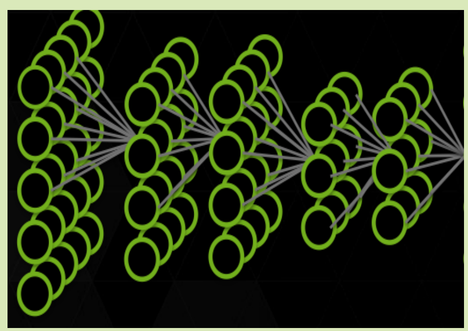

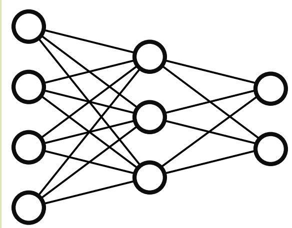

Kontainerschiff

Tiger

Zürcher Hochschule
für Angewandte Wissenschaften

Grundidee Deep Learning: “feature learning”

Automatisierung komplexer Prozesse basierend auf (hoch-dimensionalem) Sensor-Input

Bildklassifikation (herkömmlicher Ansatz)

Merkmalsextraktion manuell definierter Deskriptoren

(0.2, 0.4, …)

(0.4, 0.3, …)

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

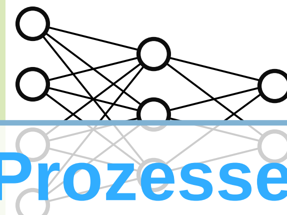

Kontainerschiff

Tiger

Bildklassifikation (neu: Convolutional Neural Networks)

Rohdaten als Input, wesentliche Merkmale werden automatisch gelernt

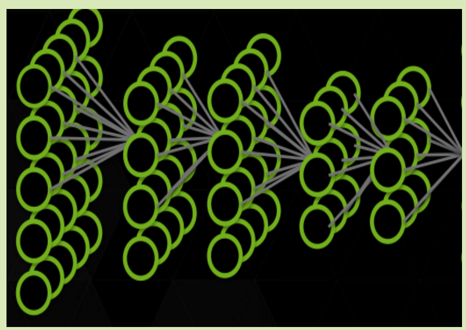

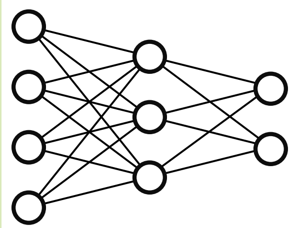

Kontainerschiff

Tiger

Prinzip: Suche der Parameter einer Funktion

Neuron

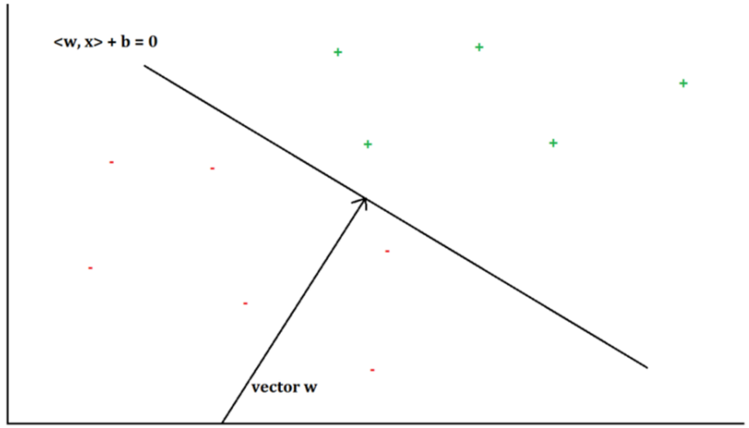

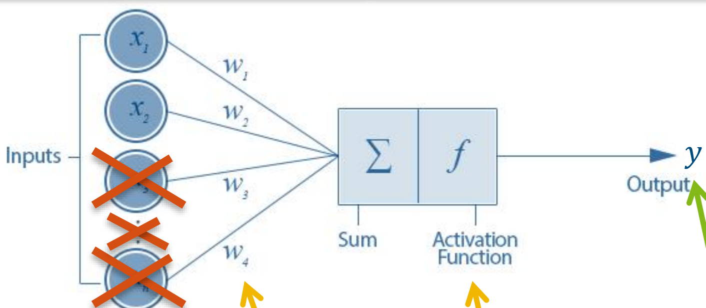

Entscheidung (Schwellwert)

Merkmale (z.B. Pixel)

Ergebnis (z.B. «Auto»)

Anpassbare Parameter

Neuronales Netz

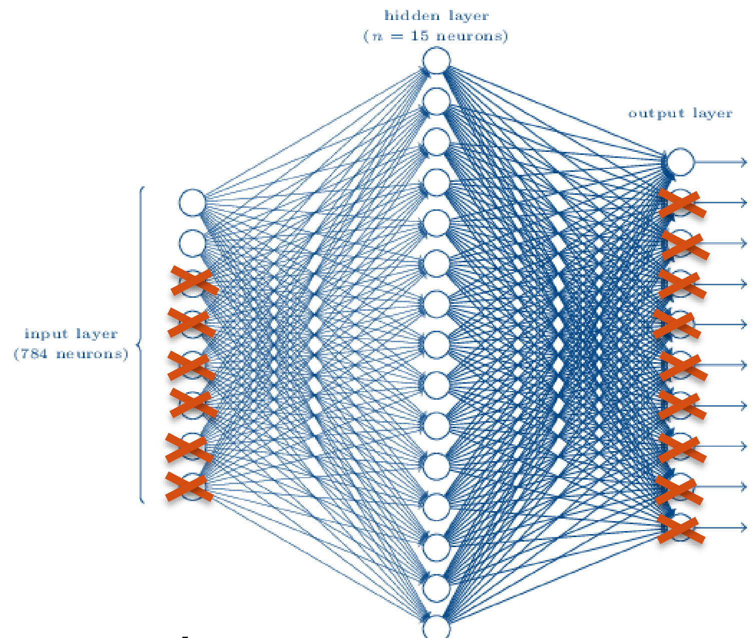

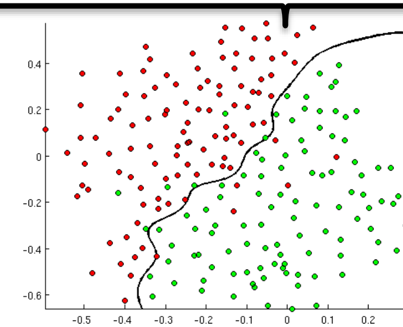

Was→ Wo?→ Wohin?

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]
  - [ML-Heavy] T
  - Completing y
- Conclusion
- Partial bibliography
- Bonus: Incomplete

Introduction

Content-aware fill is a po completion and inpaintin 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

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

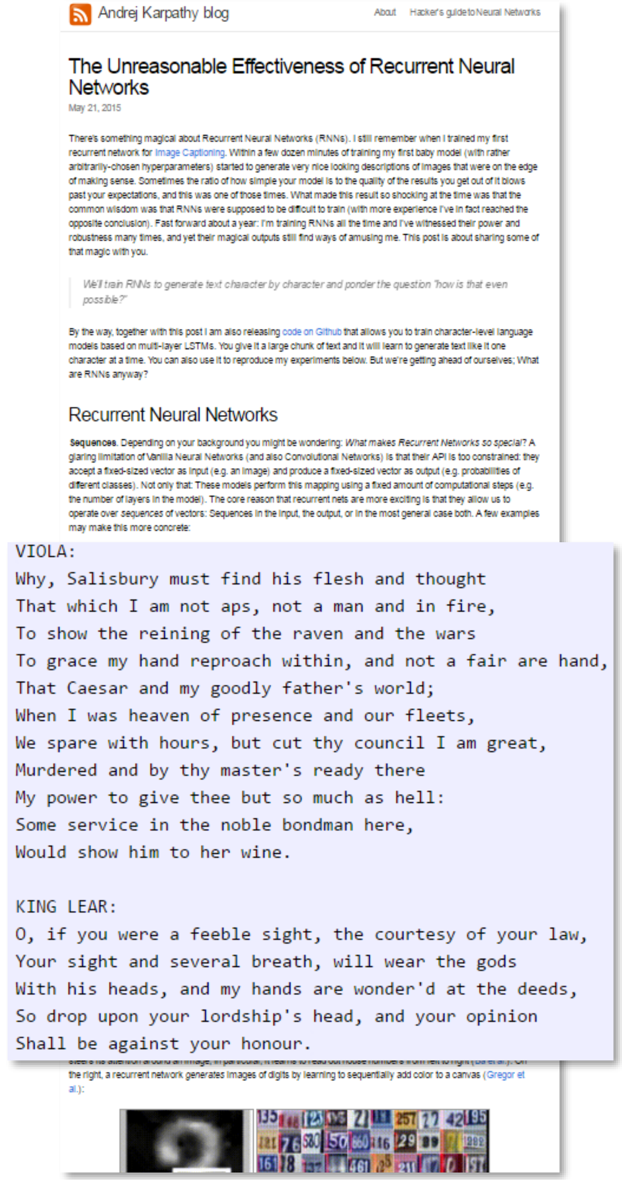

Andrej Karpathy blog   About   Hacker's guide to Neural Networks

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.):

Einige Beispiele aus den Schlagzeilen

Zürcher Hochschule
für Angewandte Wissenschaften

zh
aw

Image Completion with Deep Learning in TensorFlow

August 9, 2016

Introduction

Content-aware fill is a p
completion and inpaintin
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

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.):

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

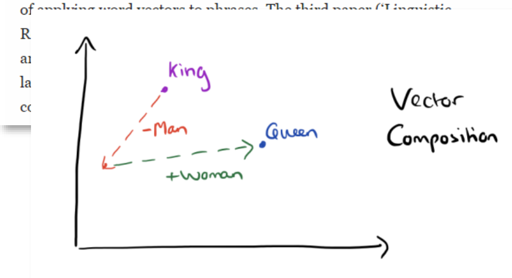

Einige Beispiele aus den Schlagzeilen

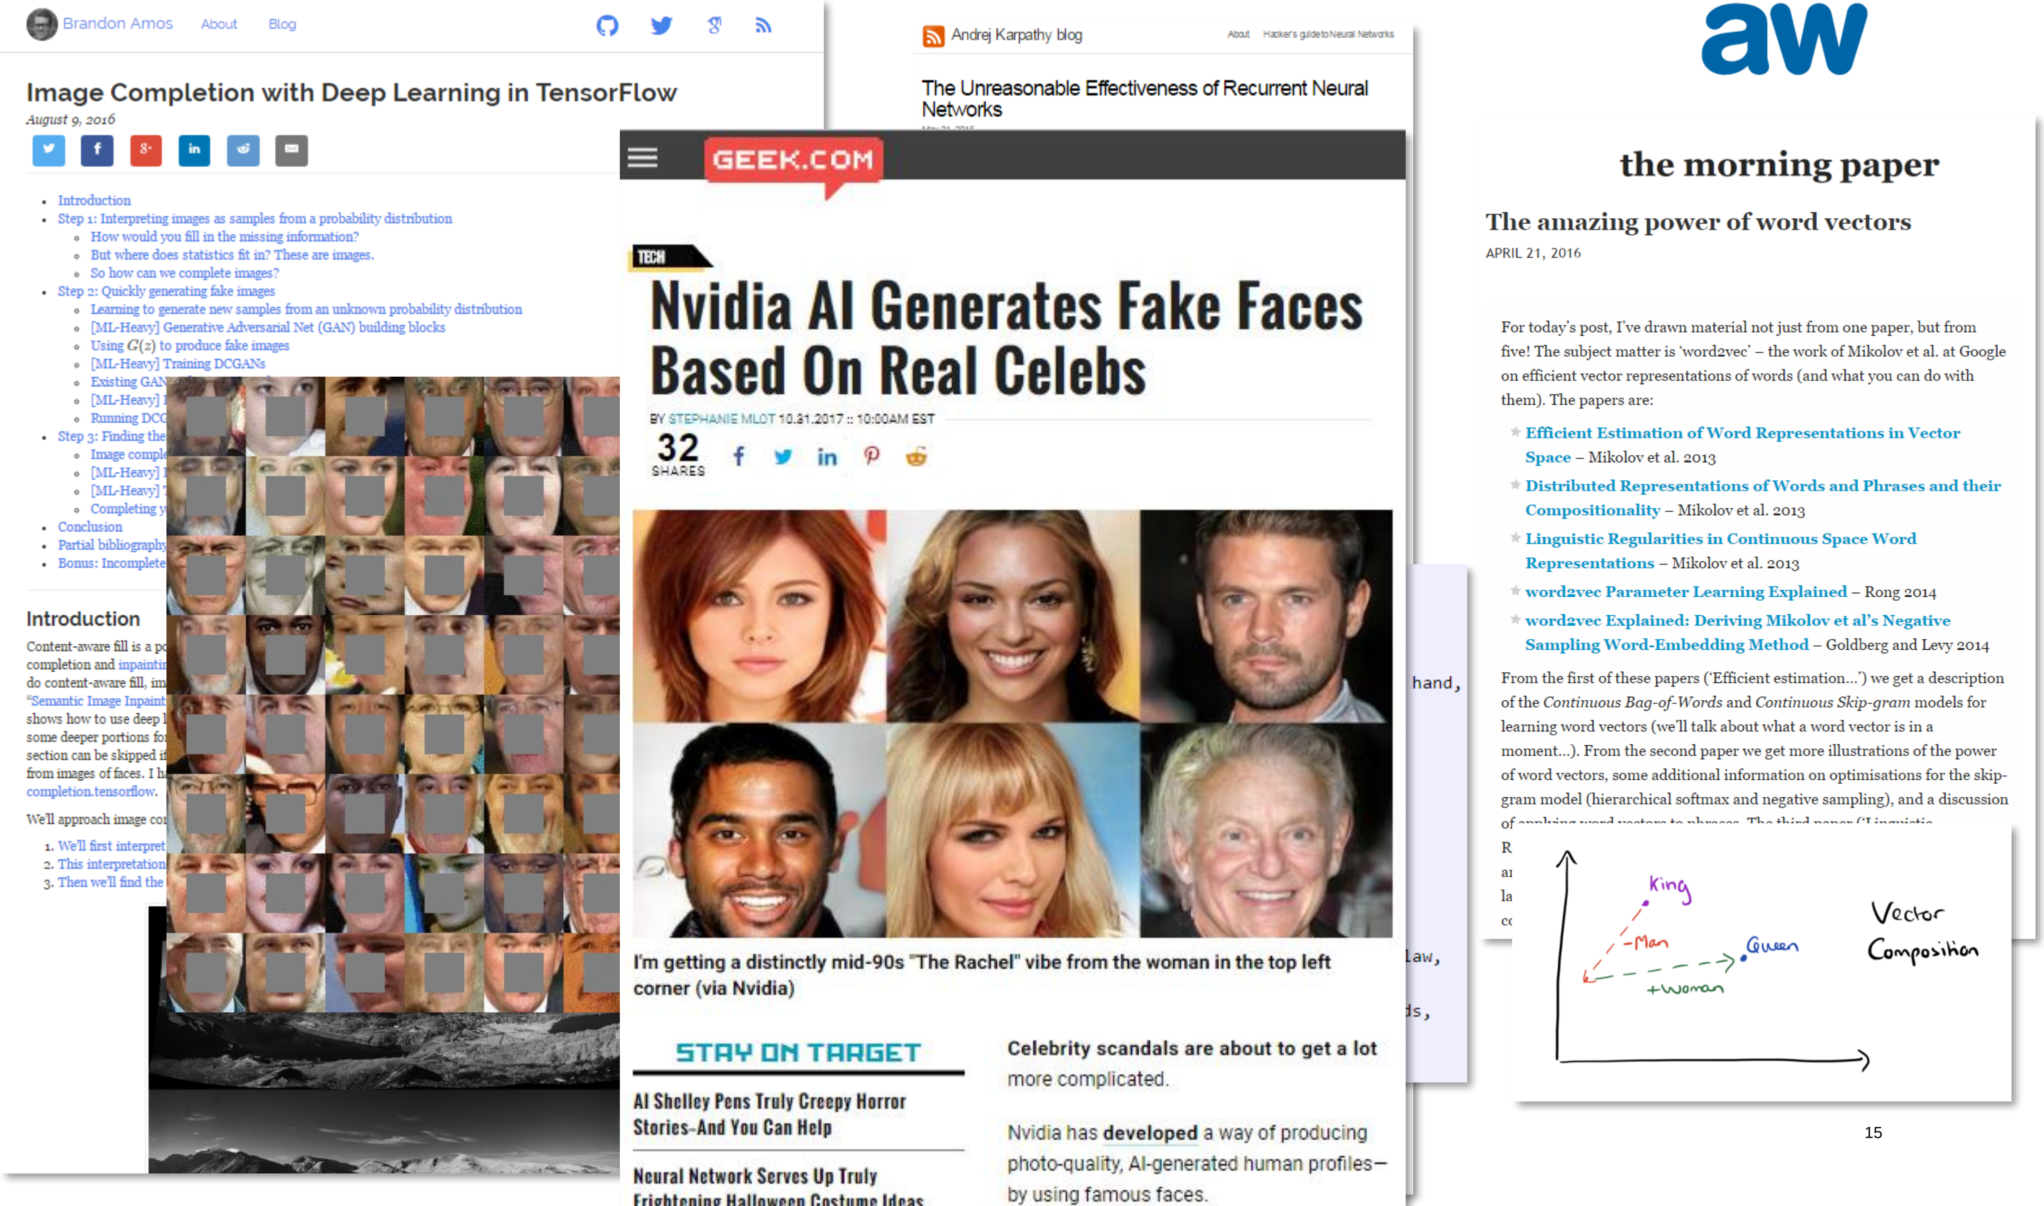

Marktchancen: Machbarkeit vs. Verantwortung
Technologie: Computer Vision mit Deep Learning

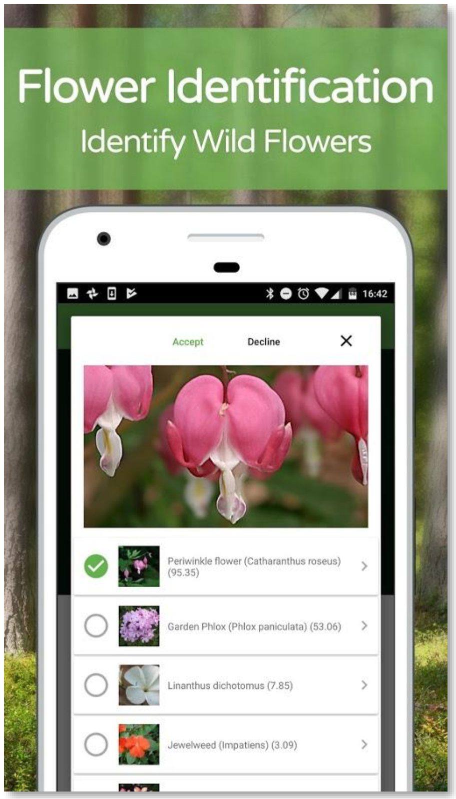

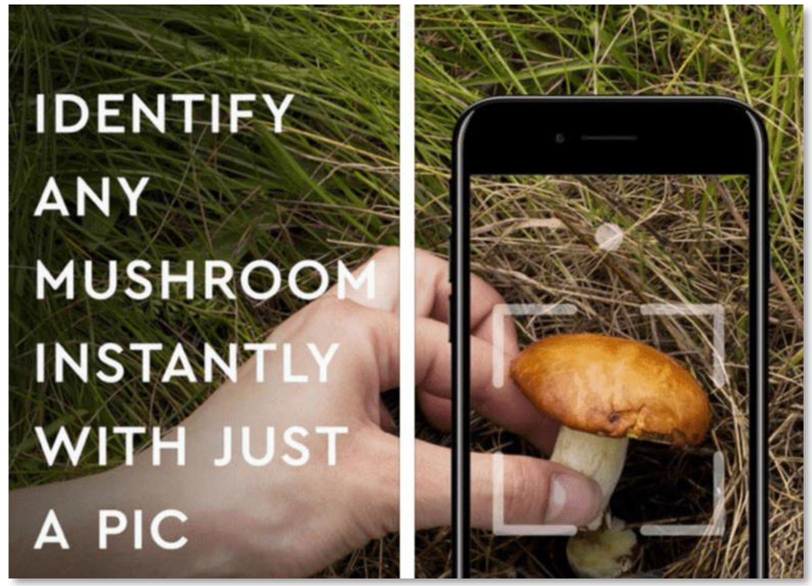

https://www.cultofmac.com/495088/avoid-potentially-deadly-ai-app/

Marktchancen: Markterfolg vs. Regulierung
Technologie: Recommender System

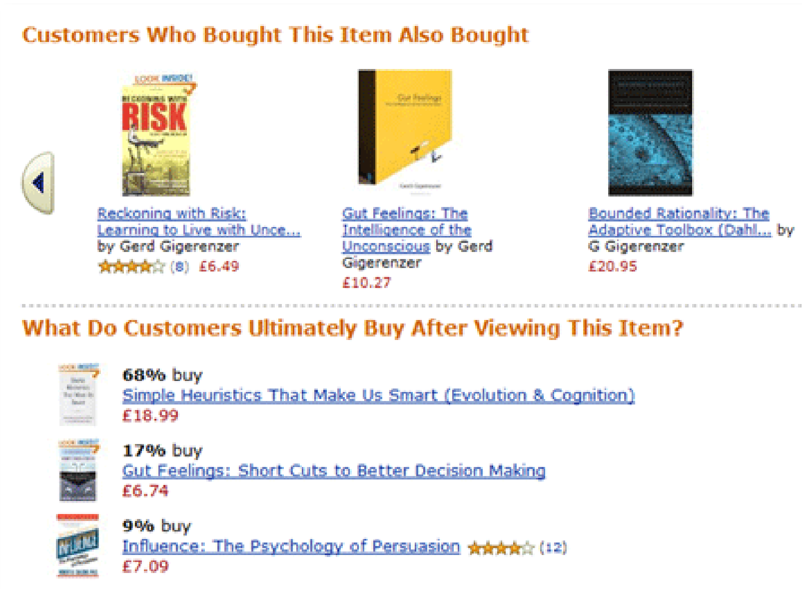

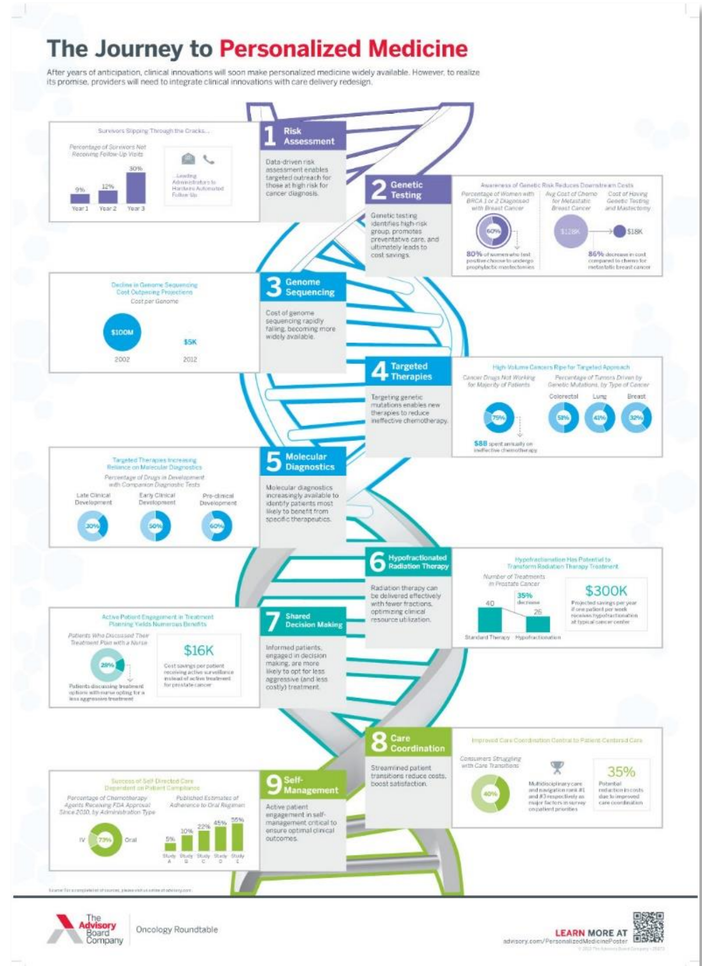

Marktchancen: Statistik vs. Bias
Technologie: Machine Learning

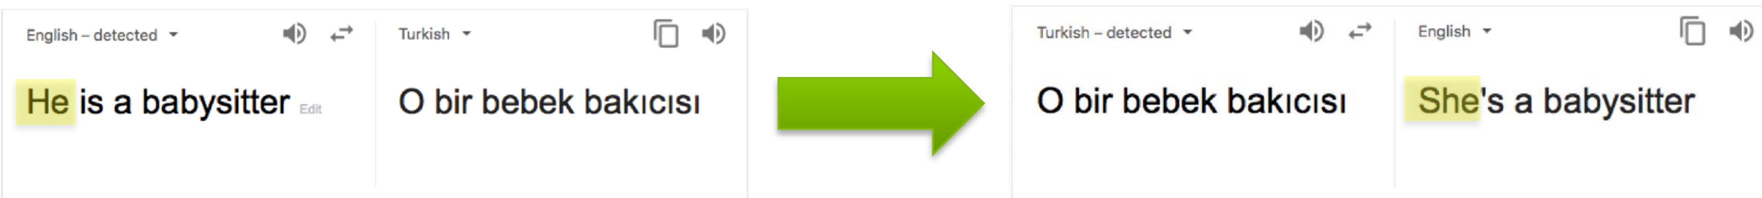

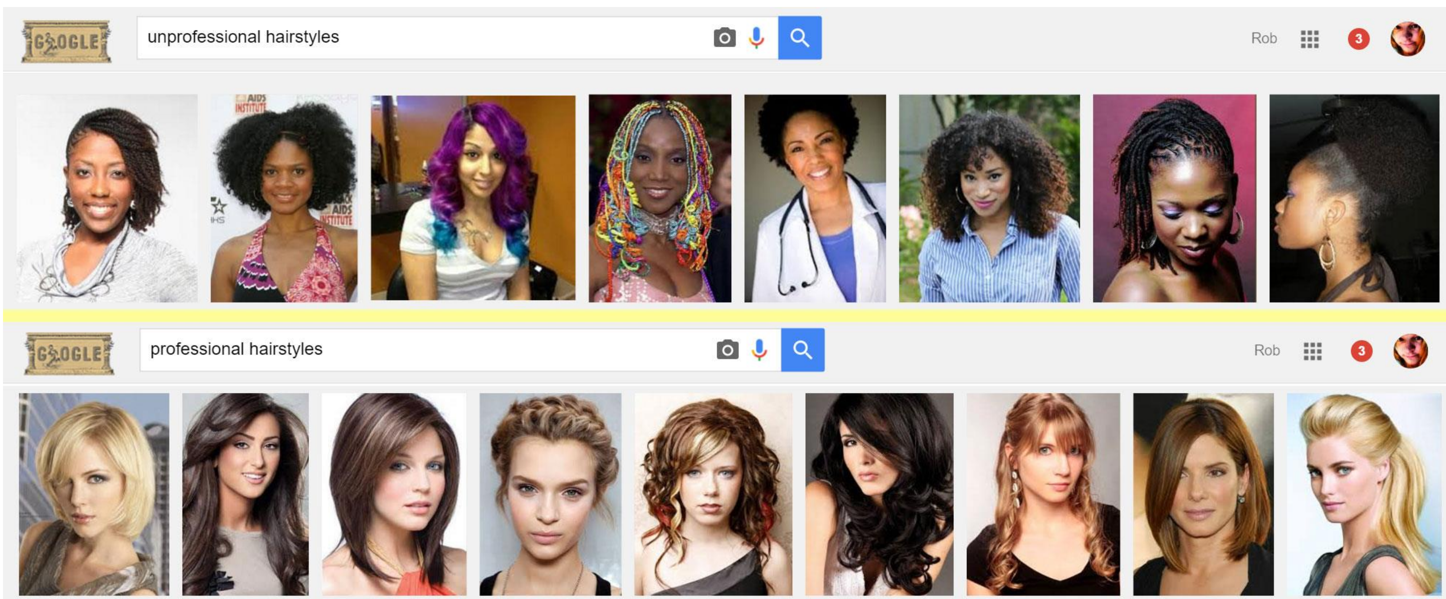

See also: Nassim Nicholas Talib, «*The Black Swan: The Impact of the Highly Improbable*», 2007

Zurich University
of Applied Sciences

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

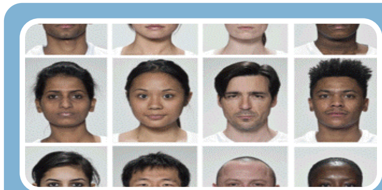

Gesichtserkennung für Stadionzutritt

[!] DEEPIMPACT

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

Vermögen beschlagnahmt

Asylbewerber können bleiben

Klage von Le Pen abgelehnt

Automatische Artikelsegmentierung

ARGUS DATA INSIGHTS

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

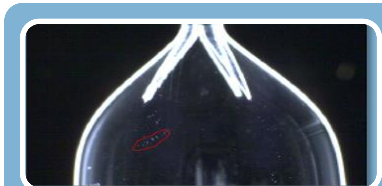

Visuelle Qualitätskontrolle in Produktion

BW-TEC

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

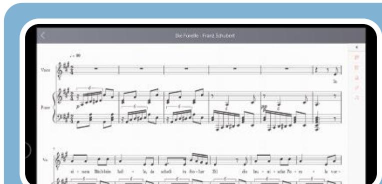

Digitalisierung von Musikalien

SCOREPAD

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

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:

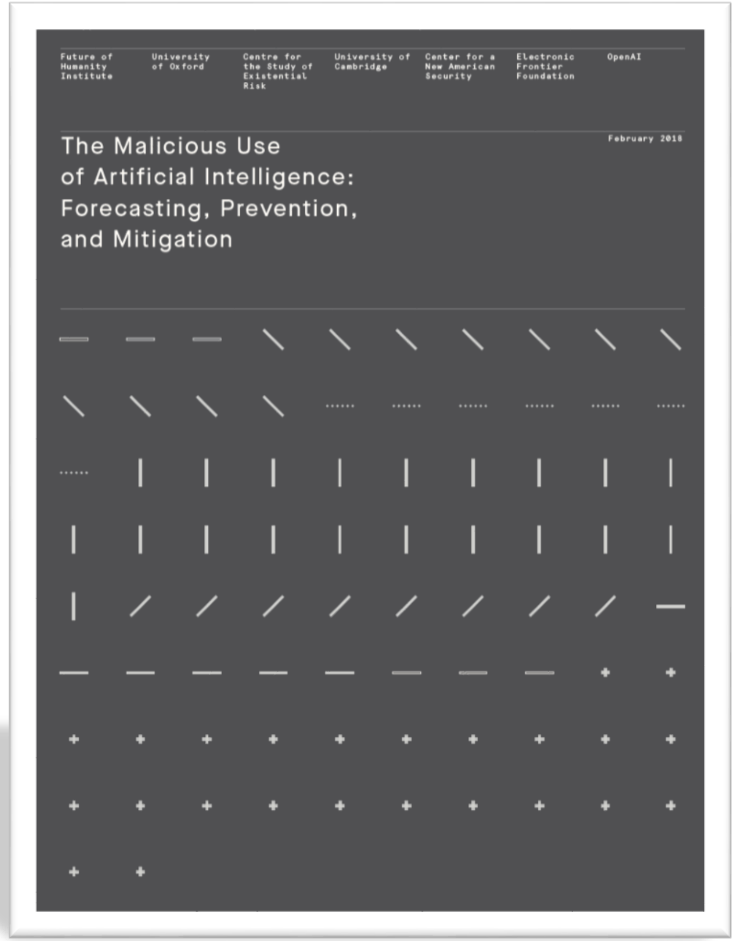

Was→ Wo?→ Wohin?

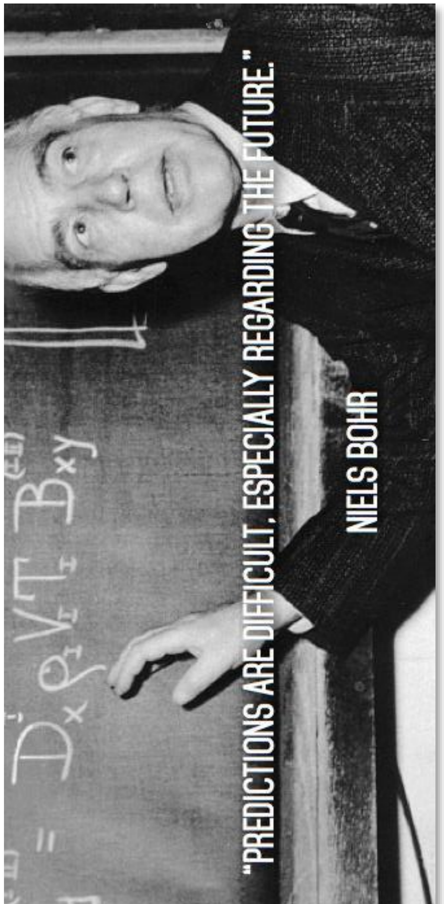

3

Wohin mag das führen?

Zurich University
of Applied Sciences

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

KI

Enorm erweiterte Automatisierungstiefe durch Fortschritt in Mustererkennung

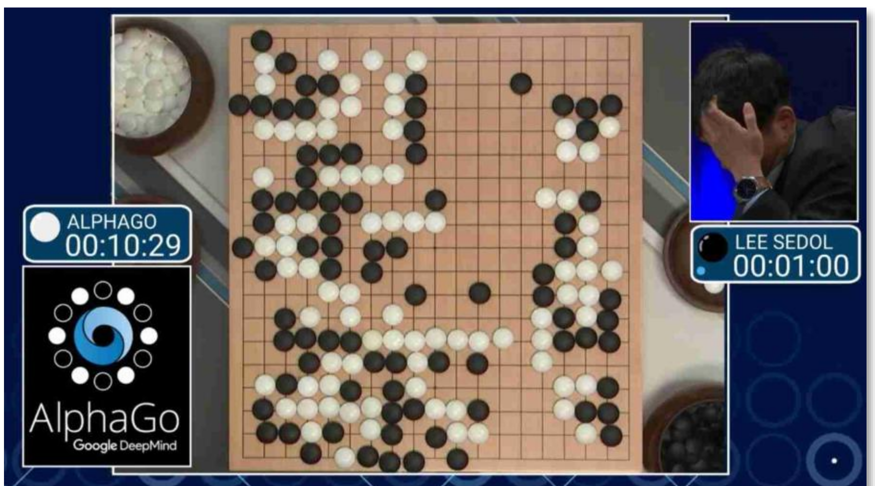

CLOUD COMPUTING

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

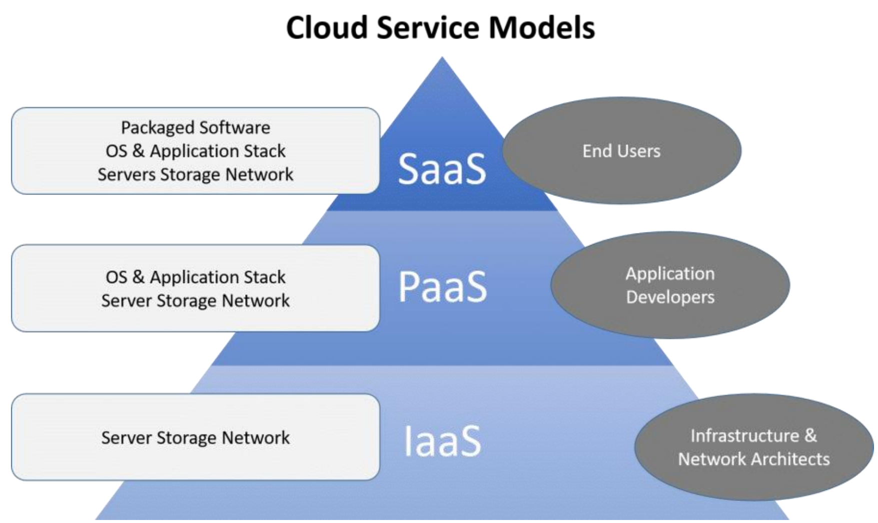

Zurich University of Applied Sciences

Grundlagen des disruptiven Potentials (II): Entkopplung

Grösse der Idee ≠ Grösse des Unternehmens

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

Technologie ist branchenunabhängig

…was **neue** Kooperationen und Allianzen ermöglicht

Grundlagen des disruptiven Potentials (III): Geschwindigkeit

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

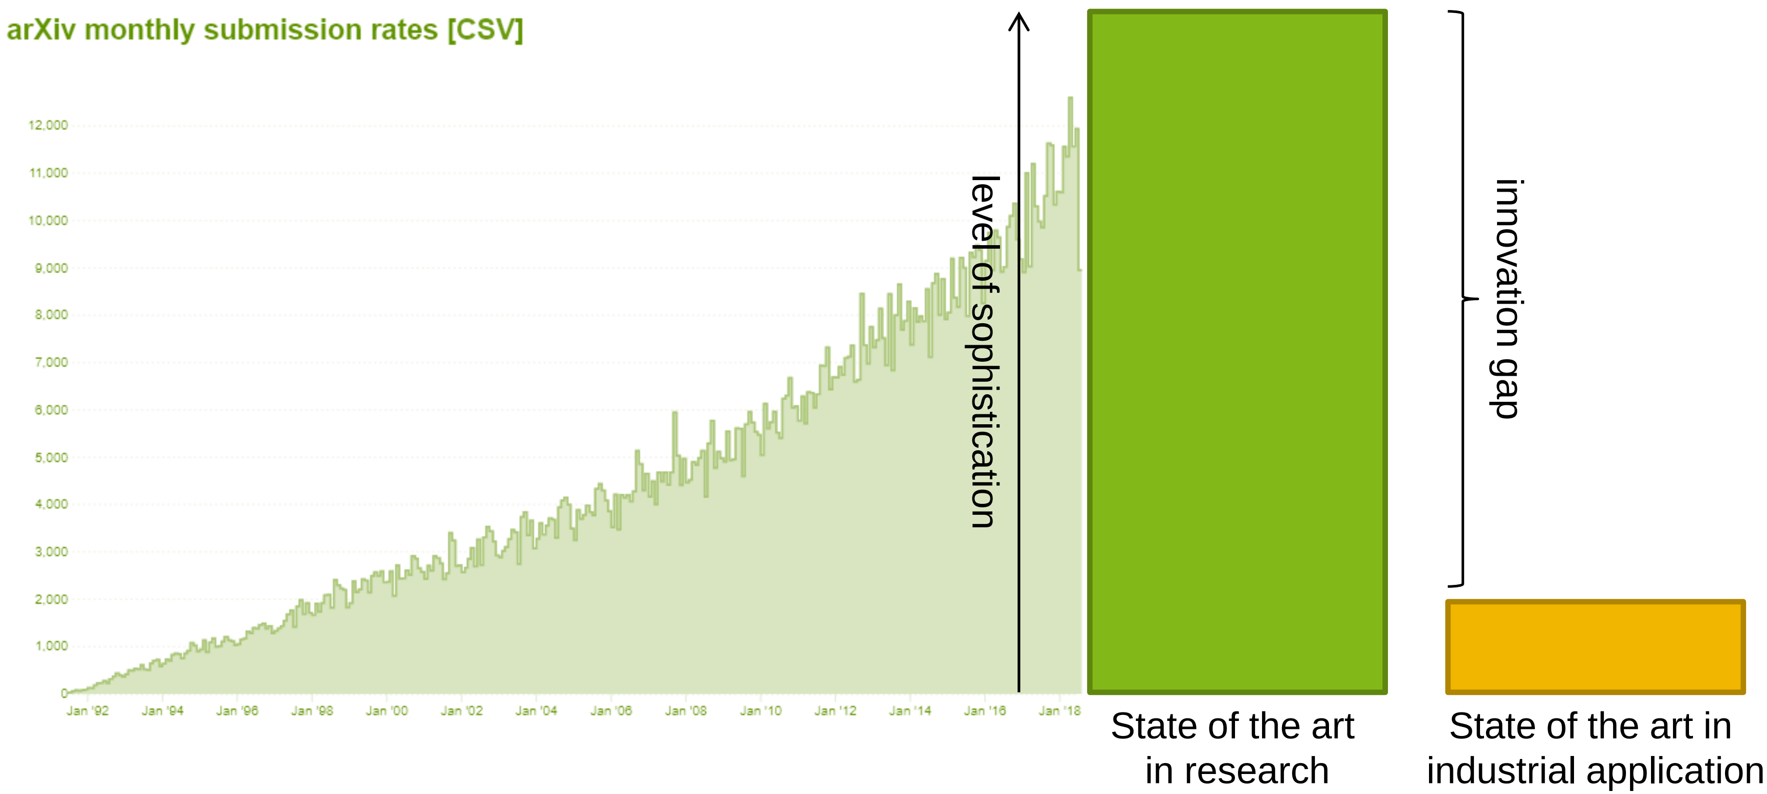

Aussicht: Disruption
…selbst bei völliger Stagnation des technischen Fortschritts

1. Hypothese: Einsatz (aktueller) KI wird sich massiv ausbreiten (Zeitrahmen: 5 Jahre)
- Indikator: **KI-Fortschritt** 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.

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”)

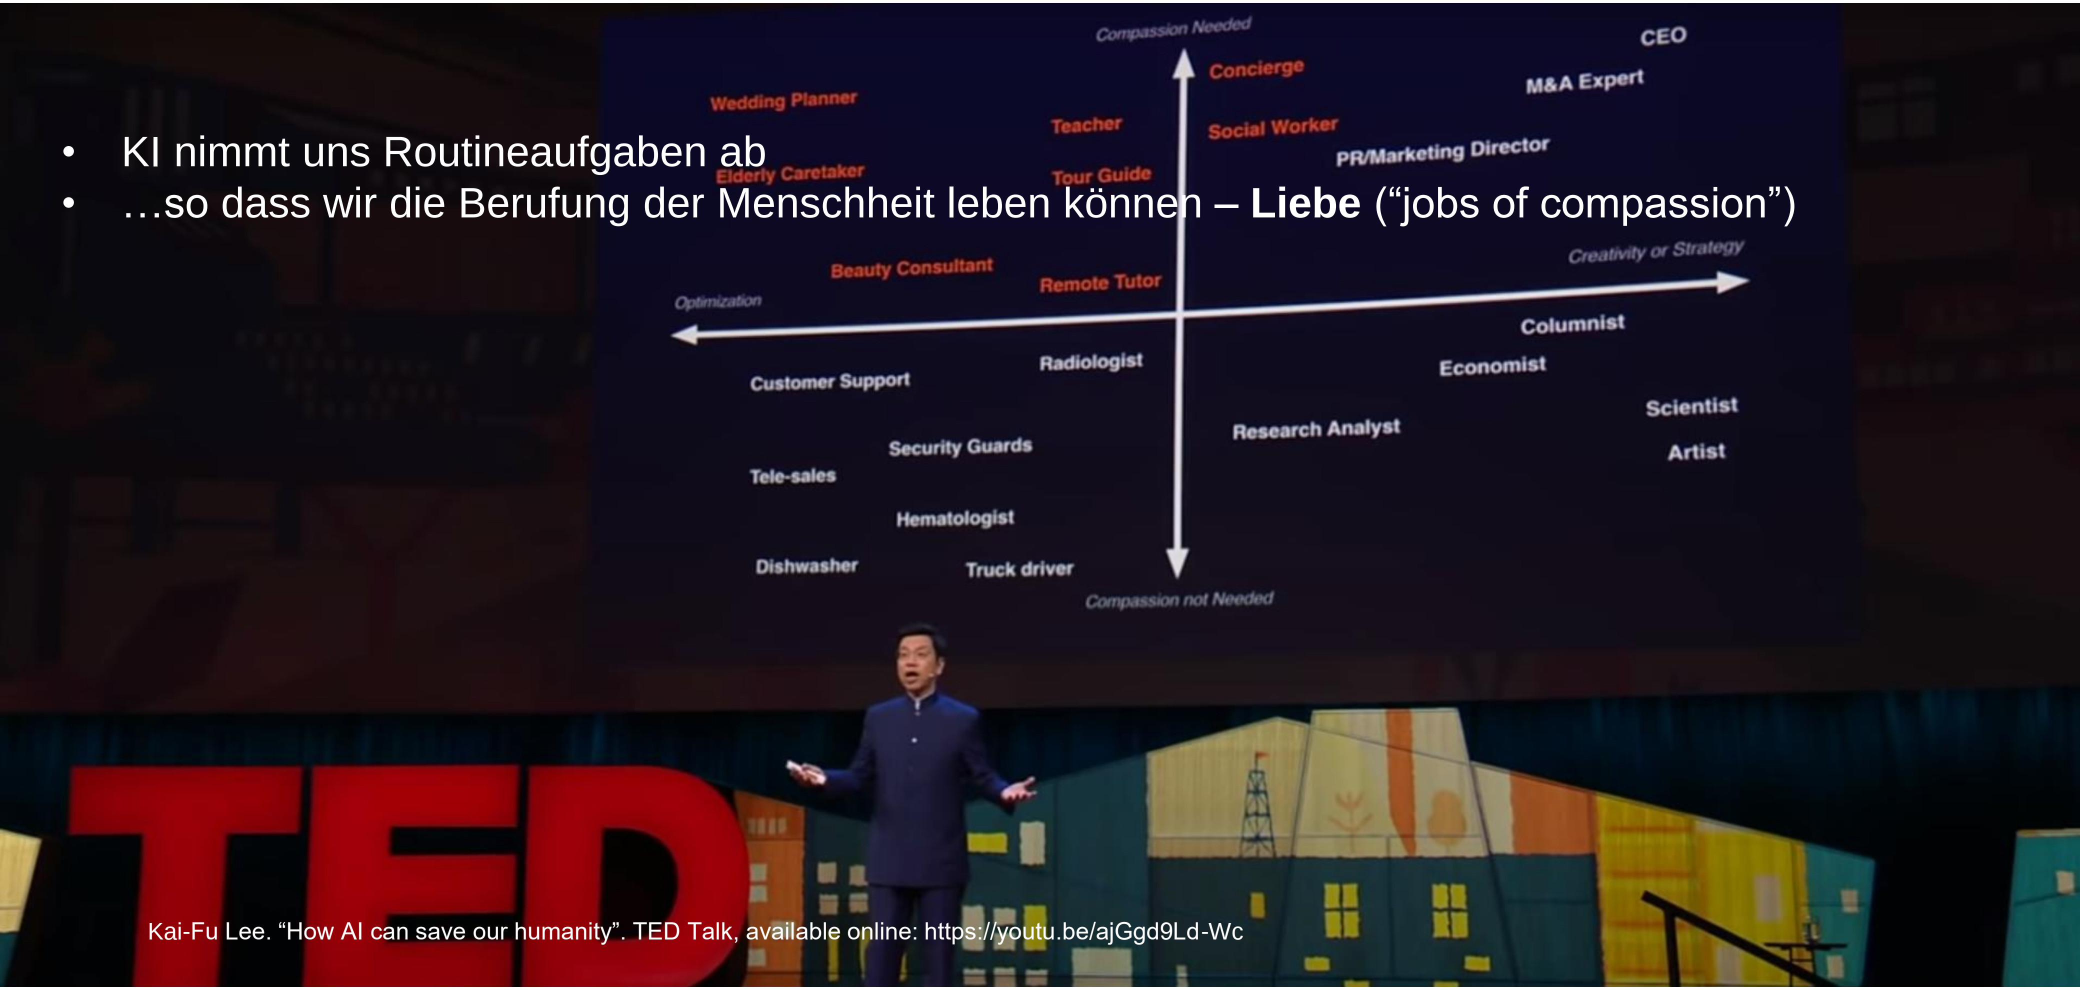

Kai-Fu Lee. “How AI can save our humanity”. TED Talk, available online: https://youtu.be/ajGgd9Ld-Wc

Zusammenfassung

- **Paradigmenwechsel** durch 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)

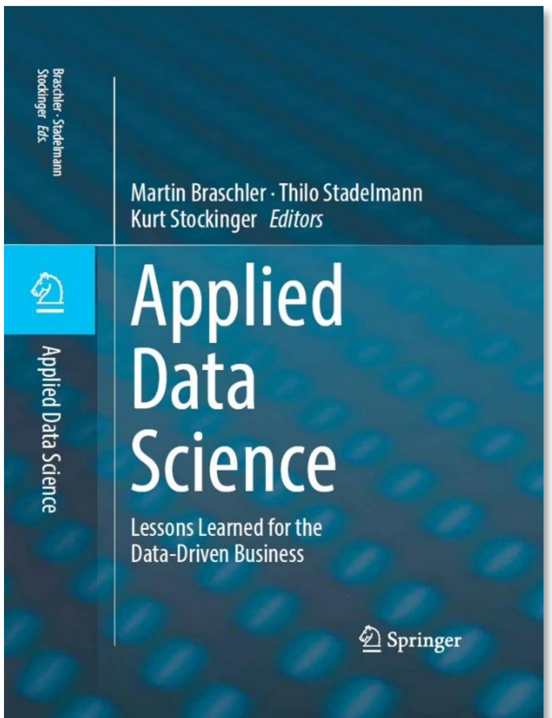

Über mich:
- Prof. AI/ML, scientific director ZHAW digital, board Data+Service
- Emai: thilo.stadelmann@zhaw.ch
- Phone: +41 58 934 72 08
- Twitter, LinkedIn: @thilo_on_data, thilo-stadelmann
- Web: https://stdm.github.io/

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

➔ Ich freue mich auf die Diskussion und generell über Kontakt!

datalab www.zhaw.ch/datalab

data service Swiss Alliance for Data-Intensive Services