

Künstliche Intelligenz – was, wo & wohin?

Volkshochschule Winterthur, 31. Oktober 2019

Thilo Stadelmann

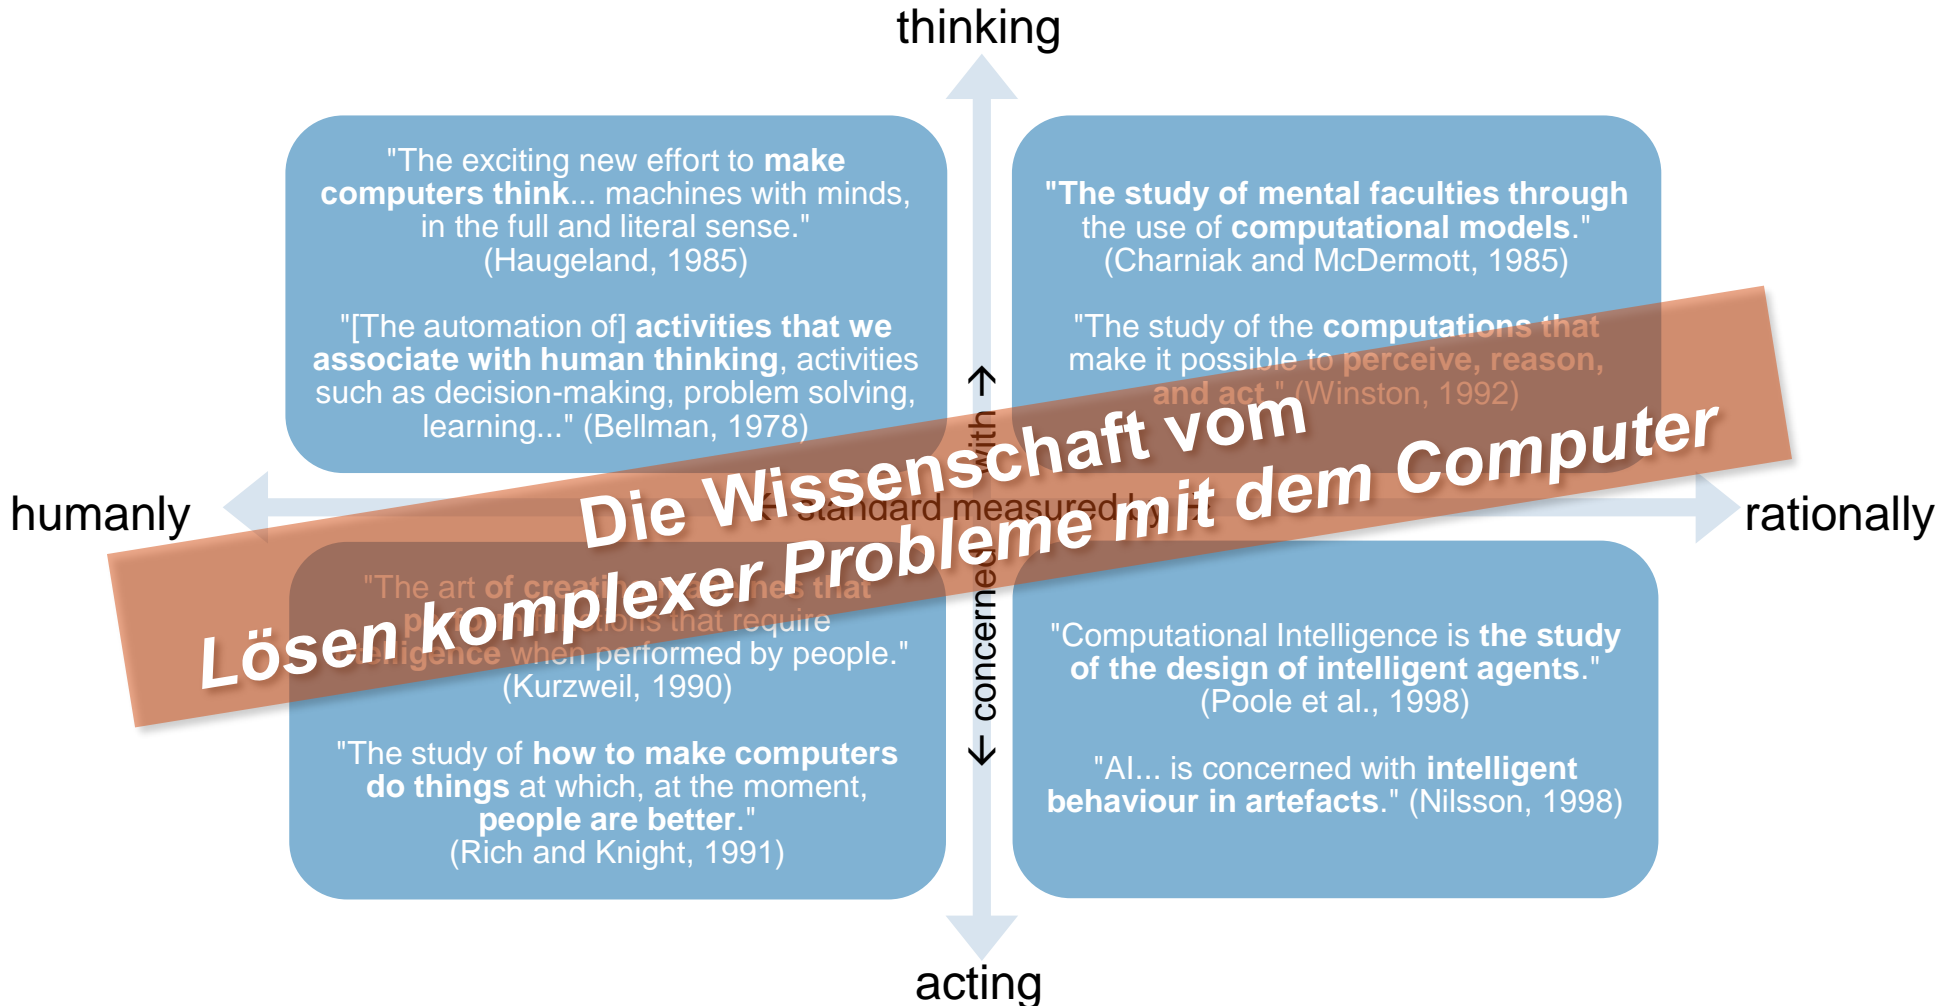


Was → Wo? → Wohin?

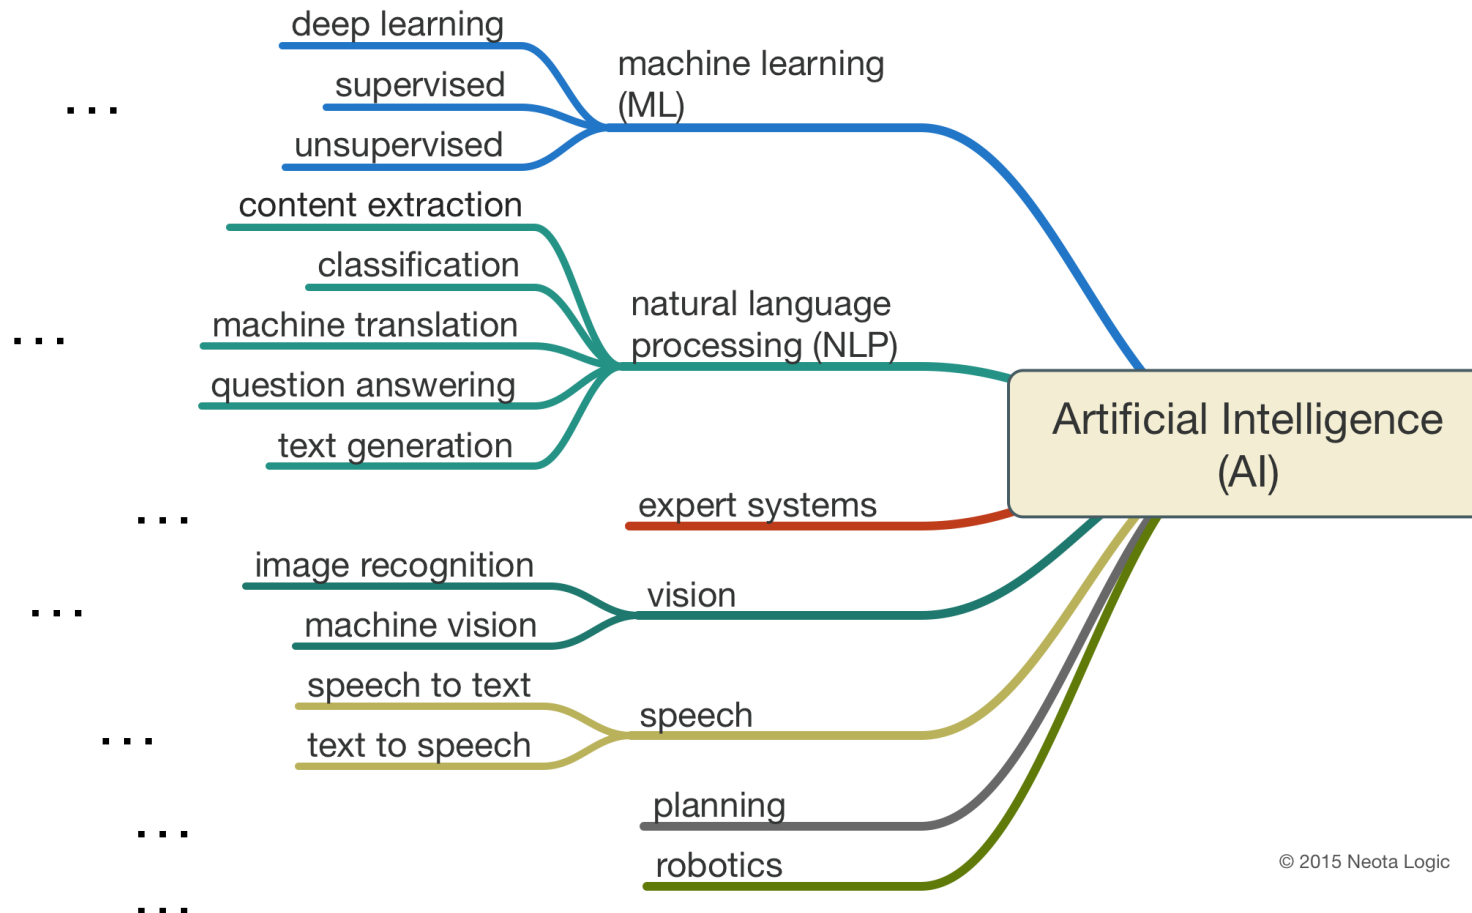
1

Was ist Künstliche Intelligenz?

Was ist künstliche Intelligenz?


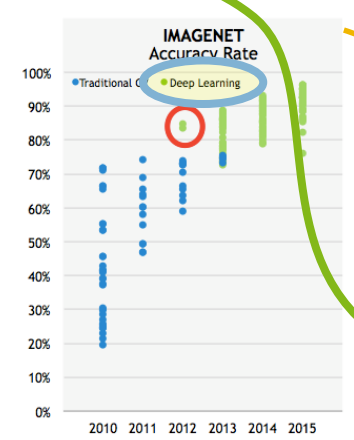




Was gehört zu künstlicher Intelligenz?



© 2015 Neota Logic

KI im Kontext

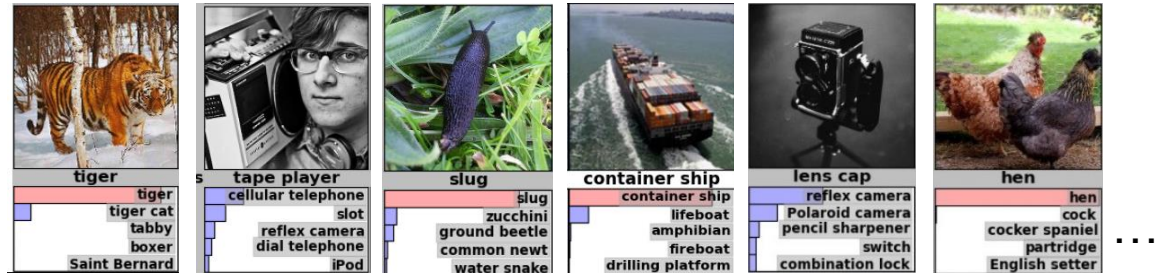


Was ist passiert?

Der ImageNet Wettbewerb



1000 Kategorien
1 Mio. Beispiele



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)

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

Idee: Mehr «Tiefe» um Merkmale automatisch zu lernen

Classical image processing




Feature extraction
(SIFT, SURF, LBP, HOG, etc.)

(0.2, 0.4, ...)

(0.4, 0.3, ...)

Classification
(SVM, neural network, etc.)



Container ship

Tiger

...

Idee: Mehr «Tiefe» um Merkmale automatisch zu lernen

Classical image processing




Feature extraction
(SIFT, SURF, LBP, HOG, etc.)

(0.2, 0.4, ...)

(0.4, 0.3, ...)

Classification
(SVM, neural network, etc.)



Container ship



Tiger

...

Using Convolutional
Neural Networks
(CNNs)



Takes raw pixels in, learns
features automatically!



Container ship

Tiger

...

Idee: Mehr «Tiefe» um Merkmale automatisch zu lernen

Classical image processing



Feature extraction
(SIFT, SURF, LBP, HOG, etc.)

(0.2, 0.4, ...)

Classification
(SVM, neural network, etc.)





Container ship

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

Using Convolutional Neural Networks (CNNs)



Takes raw pixels in, learns features automatically!




Container ship

Tiger

...


Supervised Machine Learning im Überblick

Trainingsdaten, repräsentiert durch einen Merkmalsvektor x




Supervised Machine Learning im Überblick

Trainingsdaten, repräsentiert durch einen Merkmalsvektor x



Supervised Machine Learning im Überblick


Trainingsdaten, repräsentiert durch einen Merkmalsvektor x



Wir hoffen auf gute Generalisierung bzgl. neuer Testdaten

Supervised Machine Learning im Überblick


Trainingsdaten, repräsentiert durch einen Merkmalsvektor x



Supervised Machine Learning im Überblick

Trainingsdaten, repräsentiert durch einen Merkmalsvektor x

Dieses Modell scheint gut zu passen



$$\arg \min_{h \in \mathcal{H}} \sum_{(x,y) \in D} \ell(y, h(x))$$

Wir hoffen auf gute Generalisierung bzgl. neuer Testdaten

Wir suchen gute Modelle in einem Hypothesenraum \mathcal{H} durch Minimierung des Loss ℓ zwischen Label y und Resultat $h(x)$

Was → Wo? → Wohin?

2

Wo wird das bereits eingesetzt?

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


Zürcher Hochschule für Angewandte Wissenschaften



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



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



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

At last — a computer program that can beat a champion Go player PAGE 484

ALL SYSTEMS GO

CONSERVATION: SONGBIRDS A LA CARTE. Illegal harvest of millions of Mediterranean birds. PAGE 452

RESEARCH ETHICS: SAFEGUARD TRANSPARENCY. Don't let openness backfire on individuals. PAGE 459

POPULAR SCIENCE: WHEN GENES GOT 'SELFISH'. Dawkins's calling card forty years on. PAGE 462

NATURE.COM/NATURE
28 January 2014 £10
Vol 529, No 7587



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

And the results are impressive

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

SHARE
 TWEET
 LINKEDIN

Computing

Algorithm Artistic Other In

A deep neural n
other images.

by Emerging Tect

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



Original photo

Reference photo

Result

Ad closed by Google


[Report this ad](#)

[AdChoices](#)



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



WaveNet lässt Computersprache natürlich klingen

von Henning Steier / 12.9.2018

Die Google-Tochter DeepMind macht auch Musik.




DeepMind lässt WaveNet Spr...

Die Google-Tochter DeepMind hat ein Spiel «Go» Schlagzeilen: es ist eines der besten menschlichen Spieler. Das Londoner Unternehmen erzeugt Sprache, die sehr natürlich klingt. Im Blogbeitrag des Unternehmens wird erklärt, dass die Technologie im großen Maßstab eingesetzt wird. Man hat...


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.



"My name is Avin!"



"My name is Avin!"



nerierte Sprache
is Texteingabe»

nerierte Musik
ne Inhaltsvorgabe»



1 Second

...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 GANs
 - [ML-Heavy] Training DCGANs
 - Running DCGANs
- Step 3: Finding the right image completion
 - Image completion
 - [ML-Heavy] Training DCGANs
 - [ML-Heavy] Training DCGANs
 - Completing y
- Conclusion
- Partial bibliography
- Bonus: Incomplete

Introduction

Content-aware fill is a popular image completion and inpainting technique. In this post, we'll do content-aware fill, inpainting, and semantic image inpainting. This post shows how to use deep learning for image completion. Some deeper portions for this section can be skipped if you're not interested in images of faces. I have a video on image completion.tensorflow.org

We'll approach image completion in three steps:

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



Andrej Karpathy blog

The Unreasonable Effectiveness of Recurrent Neural Networks

March 31, 2015



TECH

Nvidia AI Generates Fake Faces Based On Real Celebs

BY STEPHANIE MLDT 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.

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

hand,


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 of applying word vectors to phrases. The third paper ('Linguistic

R
a
l
a
c



Law,
is,

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



Gesichtserkennung für Stadionzutritt

- Nutzen: *Robustes* Personenidentifikationssystem
- Wirkung: Unterstützung bei Entwicklung; Datenqualität schränkte ein




Automatische Artikelsegmentierung

- Nutzen: vollautomatisches Produkt in niedrigem Preissegment
- Wirkung: Einführung dank *Teamausbau* gelungen



Visuelle Qualitätskontrolle in Produktion

- Nutzen: vollautomatischer Triage & Bearbeitung normaler Fälle
- Wirkung: macht *Familienunternehmen* zu Technologieanbieter



Digitalisierung von Musikalien

- Nutzen: Enabler für digitales Geschäftsmodell
- Wirkung: 5 Jahre nach Start ist entwickelte Technologie *grösstes Asset*





Was → Wo? → Wohin?

3

Wohin könnte und das einmal führen?

Beispiel: Machbar vs. gefährlich

Technologie: Computer Vision mit Deep Learning



Beispiel: Markterfolg vs. regulatorische Hürden

Technologie: Recommender Systems

Customers Who Bought This Item Also Bought

Reckoning with Risk: Learning to Live with Uncertainty by Gerd Gigerenzer
★★★★☆ (8) £6.49

Gut Feelings: The Intelligence of the Unconscious by Gerd Gigerenzer
£10.27

Bounded Rationality: The Adaptive Toolbox (Dahlsrud) by G Gigerenzer
£20.95

What Do Customers Ultimately Buy After Viewing This Item?

- 68%** buy **Simple Heuristics That Make Us Smart (Evolution & Cognition)** £18.99
- 17%** buy **Gut Feelings: Short Cuts to Better Decision Making** £6.74
- 9%** buy **Influence: The Psychology of Persuasion** ★★★★★ (12) £7.09

The Journey to Personalized Medicine

After years of anticipation, clinical innovations will soon make personalized medicine widely available. However, to realize its promise, providers will need to integrate clinical innovations with care delivery redesign.




- 1 Risk Assessment**
Data driven risk assessment enables targeted outreach for those at highest risk for cancer diagnosis.
- 2 Genetic Testing**
Genetic testing identifies high-risk groups, promotes preventative care, and identifies opportunities for cost savings.
Awareness of Genetic Risk Reduces Overstated Care Costs: Avg. Cost of Chemotherapy for Breast Cancer: \$120K vs. \$130K. 80% of women with positive breast cancer have no change in breast cancer mortality.
- 3 Genome Sequencing**
Cost of genome sequencing rapidly falling, becoming more widely available.
- 4 Targeted Therapies**
Targeted genetic mutations enables new therapies to reduce ineffective chemotherapy.
High-Risk Cancers Rise for Targeted Approach: Cancer Drugs Not Working for Majority of Patients. Percentage of Tumors Driven by Genetic Mutations, by Type of Cancer: Colorectal (50%), Lung (40%), Breast (30%). \$50 saved annually on ineffective chemotherapy.
- 5 Molecular Diagnostics**
Molecular diagnostics increasingly available to identify patients most likely to benefit from specific therapeutics.
- 6 Hypofractionated Radiation Therapy**
Radiation therapy can be delivered effectively with fewer fractions, optimizing clinical resource allocation.
Hypofractionation Now Preferred to Treat Breast Cancer: Number of Treatments for Breast Cancer: Standard Therapy (40) vs. Hypofractionation (26). \$300K projected savings per year if one patient per week receives hypofractionated treatment at typical cancer center.
- 7 Shared Decision Making**
Informed patients, engaged in decision making, are more likely to opt for less aggressive (and less costly) treatment.
Patients who discuss their treatment plan with a doctor: 25% reduction in unnecessary breast and prostate cancer treatments. \$16K cost savings per patient resulting from less aggressive treatment for prostate cancer.
- 8 Care Coordination**
Streamlined patient experience, reduces costs, boosts satisfaction.
Improved Care Coordination Critical to Patient-Centered Care: Consumers Struggling with Care Transitions: 40%. Multiple care plans and overlapping care at 15 different facilities may help for maximizing care coordination: 35%. Provider: Patient: Relative to costs, 40% of patients' care coordination is inefficient.
- 9 Self-Management**
Active patient engagement in self-management critical to ensure optimal clinical outcomes.
Reasons of Self-Directed Care Dependent on Patient Compliance: Percentage of Compliance Agents Receiving FDA Approval Since 2009 by Administration Type: IV (25%), Oral (75%). Additional Examples of Adherence to Oral Agent: 5%, 10%, 20%, 30%, 40%, 50%, 60%, 70%, 80%, 90%, 95%.

Source: The concept of the journey, identified in a white paper.

The Advisory Board Company | Oncology Roundtable | LEARN MORE AT advisory.com/PersonalizedMedicine | QR code

Beispiel: Statistik vs. Bias


Technologie: Machine Learning



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

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:



Aussicht: Disruption

...selbst bei völliger Stagnation des technischen Fortschritts

1. Hypothese: Einsatz (aktueller) KI wird sich massiv ausbreiten (Zeitraumen: 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**

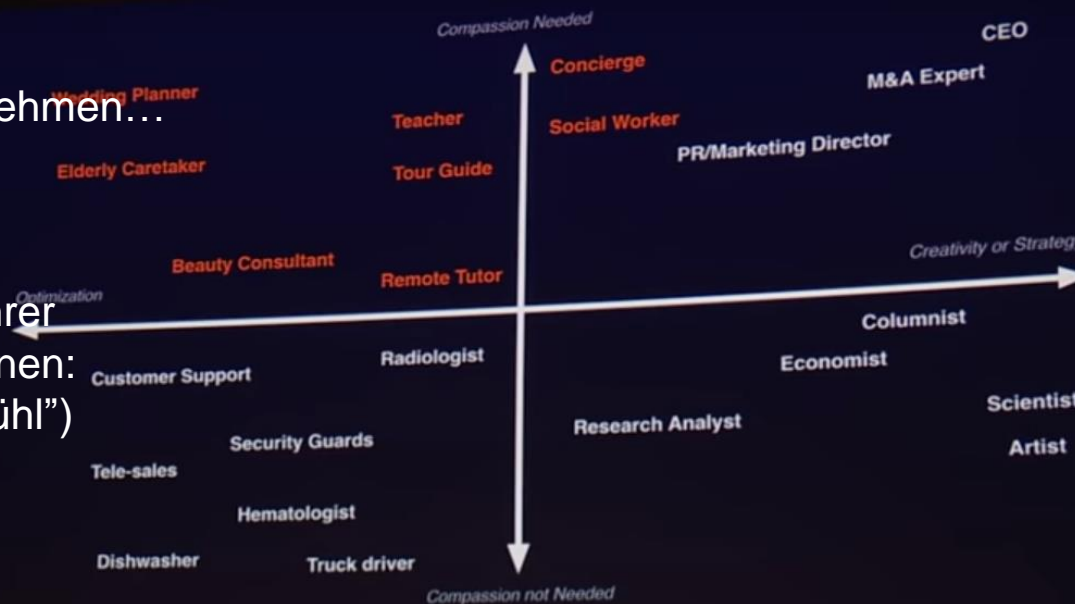


Siehe auch: Stockinger, Braschler & Stadelmann. “Lessons Learned from Challenging Data Science Case Studies”. In: Braschler et al. (Eds), “*Applied Data Science - Lessons Learned for the Data-Driven Business*”, Springer, 2019.

Eine lebenswertere Gesellschaft durch KI?

Die Vision von Kai-Fu Lee, Unternehmer & Forscher

- KI Systeme würden Routineaufgaben übernehmen...
- ...so dass **Menschen** ihrer Bestimmung folgen können: **Liebe** ("Jobs mit Mitgefühl")



Kai-Fu Lee. "How AI can save our humanity". TED Talk, available online: <https://youtu.be/ajGqd9Ld-Wc>

Schlussfolgerungen

- KI automatisiert *einzelne*, komplexe, aber *redundante* Prozesse (meist mittels maschinellem Lernen auf menschengenerierten Beispielen)
- Deep Learning hat zu Paradigmenwechsel in *Mustererkennungsaufgaben* geführt
- Das Ergebnis könnte eine menschlichere Gesellschaft sein
- Das Zeitfenster zum Gestalten beträgt wenige Jahre (<5)



Zu mir:

- Prof. KI/ML, Scientific Director ZHAW digital
- Email: stdm@zhaw.ch
- Telefon: 058 934 72 08
- Web: <https://stdm.github.io/>
- Twitter: @thilo_on_data
- LinkedIn: thilo-stadelmann



Mehr zum Thema:

- Veranstaltungen in CH: www.data-service-alliance.ch
- Lesenswert: <https://www.deeplearning.ai/thebatch/>
- Ebenso: Stuart Russell, «Human compatible», Penguin Books, 2019

ANHANG

Grundlage

Induktives überwachtetes Lernen



Annahme

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

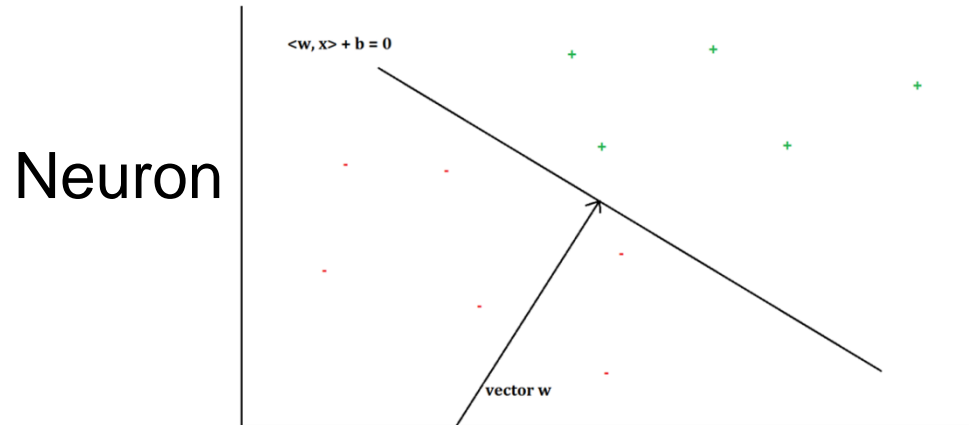
Methode

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




$$f(x) = y$$



Suche der Parameter *einer Funktion*?




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 W
($W = \{w_1, w_2, \dots\}$ 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)



$$l(W) = \frac{1}{N} \sum_{i=1}^N (f_W(x_i) - y_i)^2$$

↙ 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




Image source: "Unsupervised Learning of Hierarchical Representations with Convolutional Deep Belief Networks" ICML 2009 & Comm. ACM 2011.
Honglak Lee, Roger Grosse, Rajesh Ranganath, and Andrew Ng.

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