Zürcher Hochschule
für Angewandte Wissenschaften

Was denken denkende Maschinen?

WI-Award, Crowne Plaza Zürich, 20.10.2016

Thilo Stadelmann

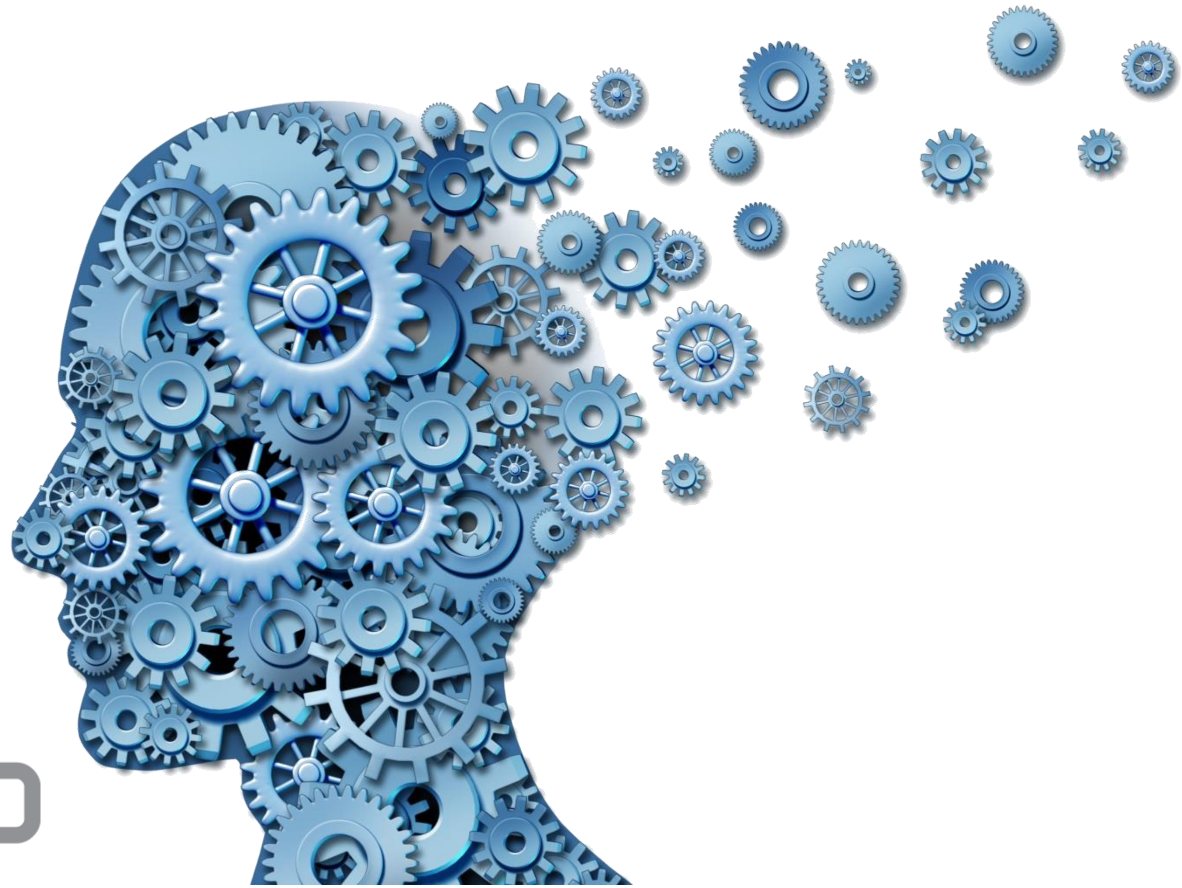

swiss group for artificial intelligence and cognitive science

Swiss Alliance for Data-Intensive Services

www.zhaw.ch/datalab

Was?→ Wie? → Wohin?

Zürcher Hochschule
für Angewandte Wissenschaften

1

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

Data Scientist:

The Sexiest Job of the 21st Century

Meet the people who can coax treasure out of messy, unstructured data.

by Thomas H. Davenport and D.J. Patil

When Jonathan Goldman arrived for work in June 2006 at LinkedIn, the business networking site, the place still felt like a start-up. The company had just under 8 million accounts, and the number was growing quickly as existing members invited their friends and colleagues to join. But users weren't seeking out connections with the people who were already on the site at the rate executives had expected. Something was apparently missing in the social experience. As one LinkedIn manager put it, "It was like arriving at a conference reception and realizing you don't know anyone. So you just stand in the corner sipping your drink—and you probably leave early."

70 Harvard Business Review **October 2012**

Data Scientist:
The Sexiest Job of the 21st Century

Meet the people who can coax treasure out of messy, unstructured data.
by Thomas H. Davenport and D.J. Patil

When Jonathan Goldman arrived for work in June 2006 at LinkedIn, the business networking site, the place still felt like a start-up. The company had just under 8 million accounts, and the number was growing quickly as existing members invited their friends and colleagues to join. But users weren't seeking out connections with the people who were already on the site at the rate executives had expected. Something was apparently missing in the social experience. As one LinkedIn manager put it, "It was like arriving at a conference reception and realizing you don't know anyone. So you just stand in the corner sipping your drink—and you probably leave early."

70 Harvard Business Review October 2012

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

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

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

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

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

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

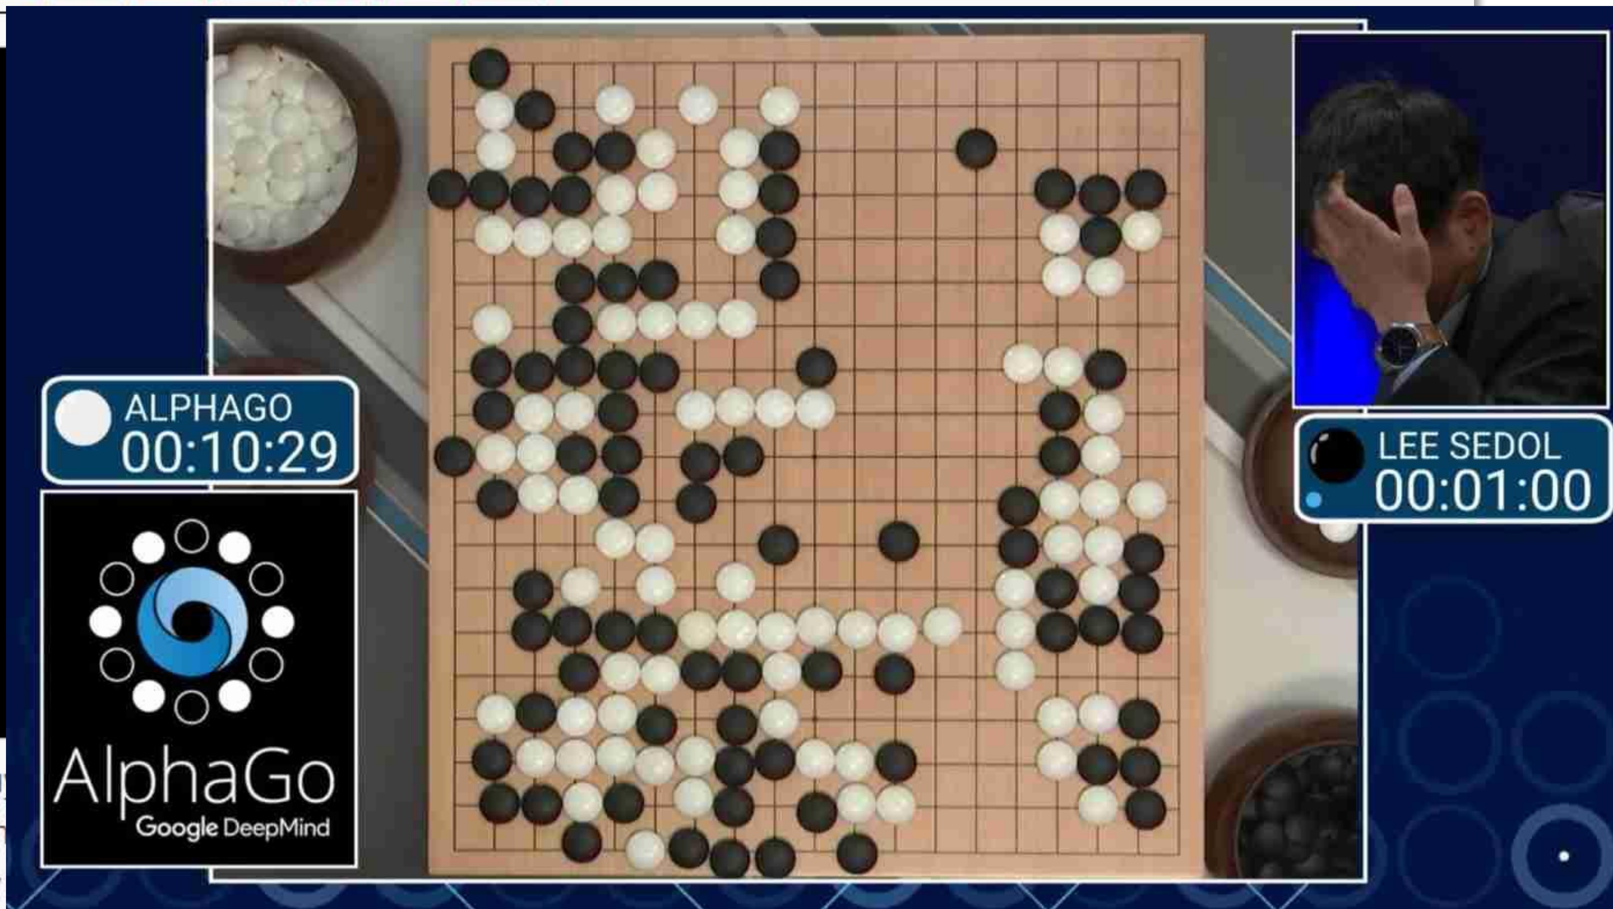

Google will bu
reports that th
in talks to buy
couldn't disclose deal terms.

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

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

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

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

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

Zürcher Hochschule
für Angewandte Wissenschaften

Künstliche Intelligenz

WaveNet lässt Computersprache natürlich klingen

von Henning Steier / 12.9.2016, 10:05 Uhr

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

KOMMENTARE

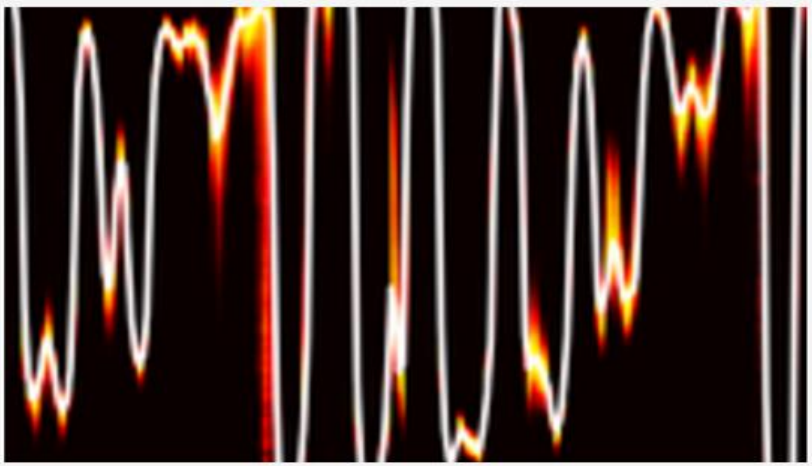

DeepMind lässt WaveNet Sprachwellen erzeugen. (Symbolbild: PD)

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

MEISTGELESEN

Künstliche Intelligenz

Kein Google für jeden

KOMMENTAR / Henning Steier / 5.10.2016

Neue Produkte aus Mountain View

Google macht sich nicht nur im Wohnzimmer breit

Henning Steier / 4.10.2016

Dropbox

68 Millionen verschlüsselte Passwörter im Netz

5.10.2016

Zürcher Hochschule für Angewandte Wissenschaften

zh aw

Generierte Sprache «aus Texteingabe»

Generierte Musik «ohne Inhaltsvorgabe»

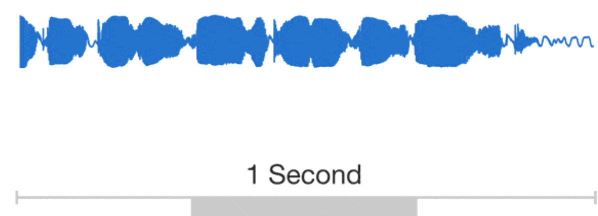

MIT Technology Review

Topics+ Top Stories

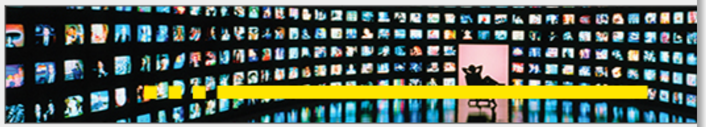

Zürcher Hochschule
für Angewandte Wissenschaften

Computing

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

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

by Emerging Technology from the arXiv May 10, 2016

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

MIT Technology Review

Topics+ Top Stories

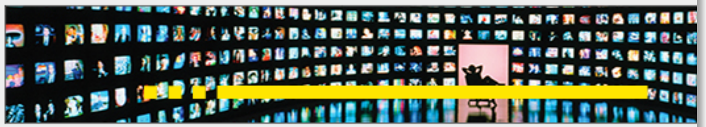

Computing

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

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

by Emerging Technology from the arXiv May 10, 2016

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

Zürcher Hochschule für Angewandte Wissenschaften

zh aw

MIT Technology Review

Topics+ Top Stories

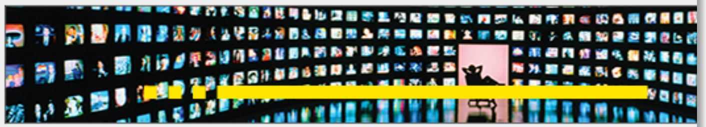

Computing

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

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

by Emerging Technology from the arXiv May 10, 2016

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

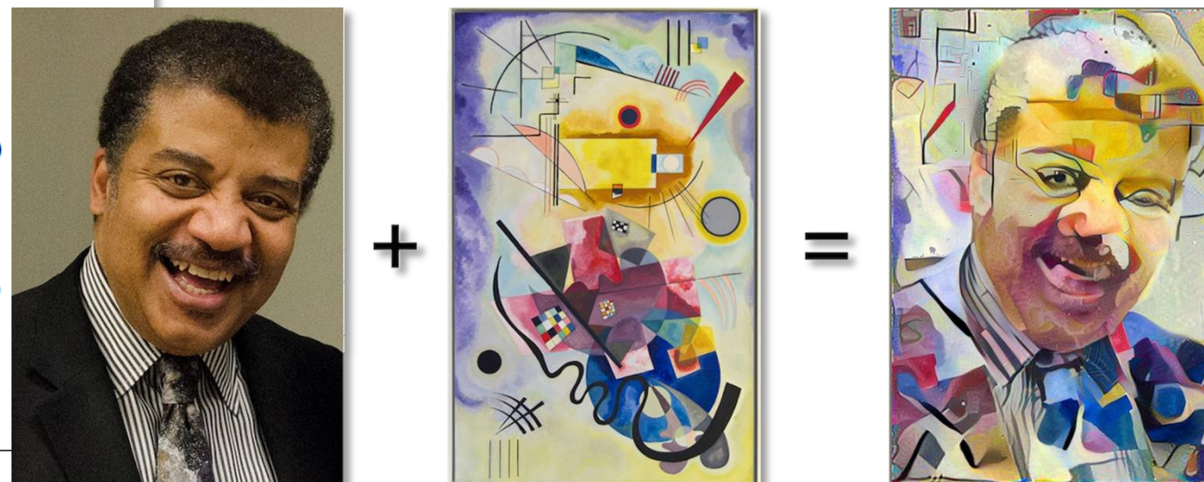

Zürcher Hochschule
für Angewandte Wissenschaften

Computing

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

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

by Emerging Technology from the arXiv May 10, 2016

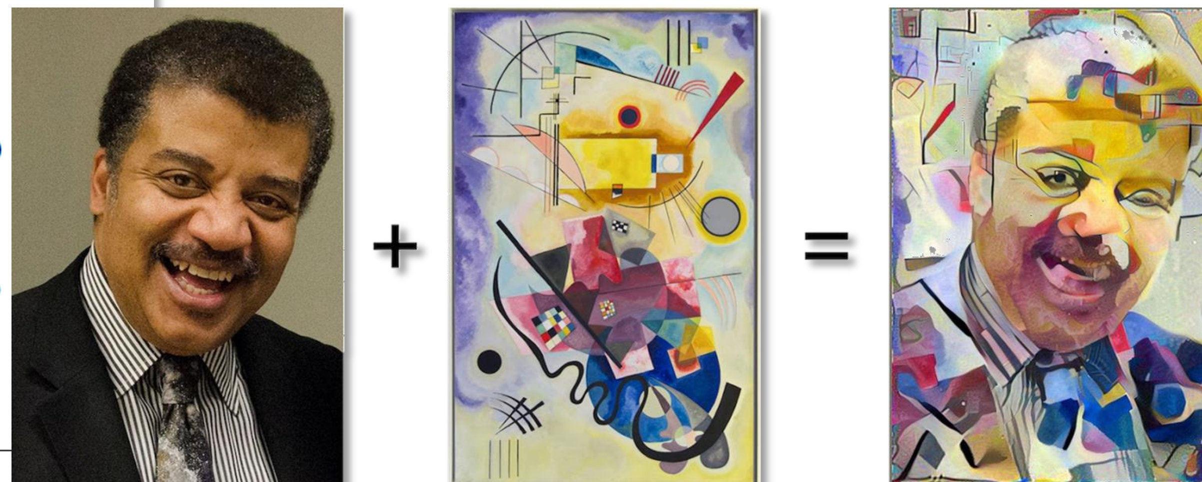

The nature of artistic style is something of a mystery to most people. Think
of Vincent Van Gogh's *Starr* n, or
Edvard Munch's *The Scream* e th
humans recognize easily.

…und die Liste liesse sich fortsetzen!

Zürcher Hochschule
für Angewandte Wissenschaften

…und die Liste liesse sich fortsetzen!

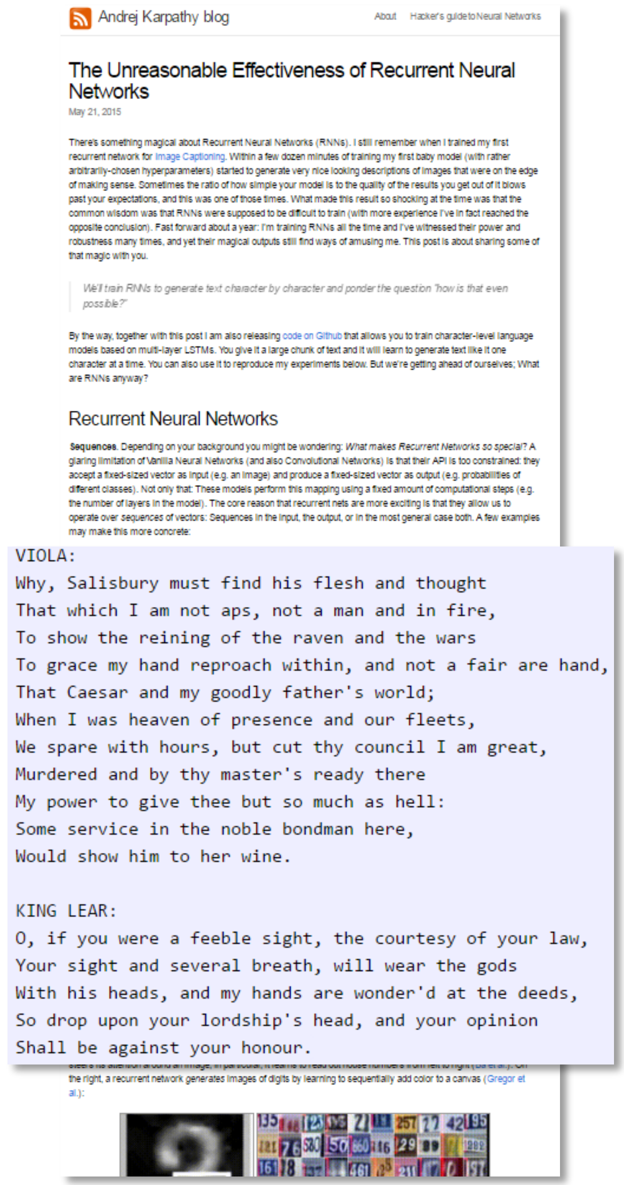

Zürcher Hochschule
für Angewandte Wissenschaften

…und die Liste liesse sich fortsetzen!

Brandon Amos About Blog

Image Completion with Deep Learning in TensorFlow

August 9, 2016

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

Introduction

Content-aware fill is a p
completion and 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

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.):

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

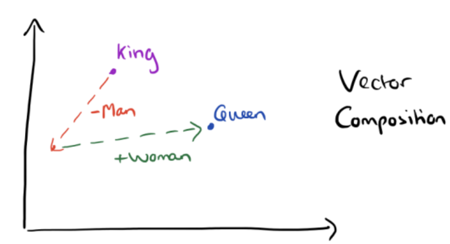

Was ist passiert?
Der ImageNet Wettbewerb

1000 Kategorien
1 Mio. Beispiele

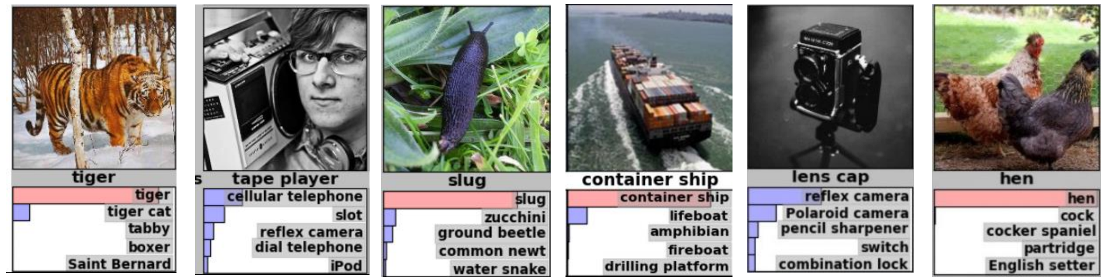

…

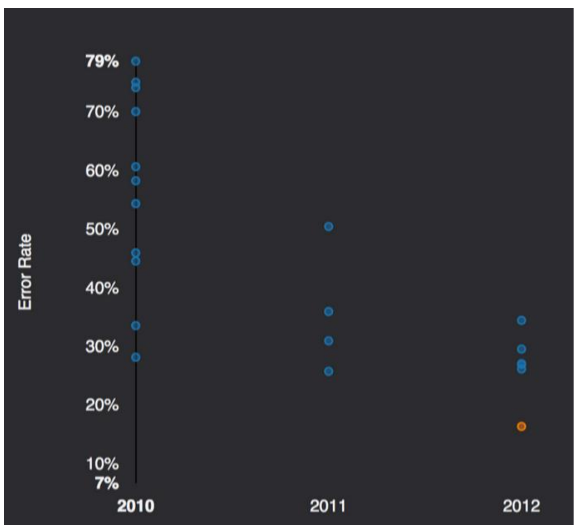

Was ist passiert?

Der ImageNet Wettbewerb

1000 Kategorien
1 Mio. Beispiele

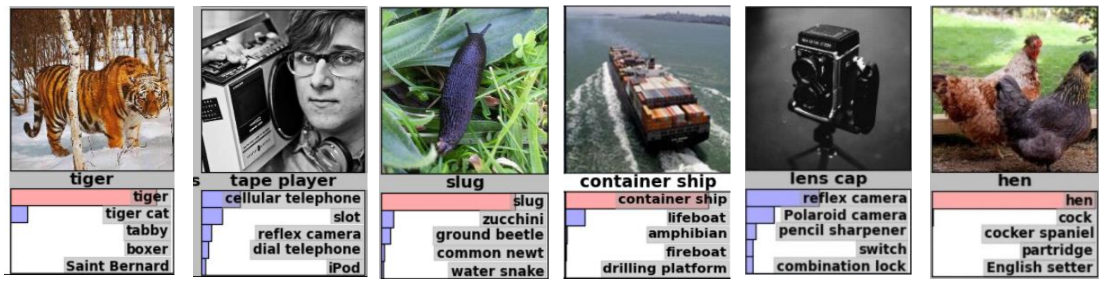

…

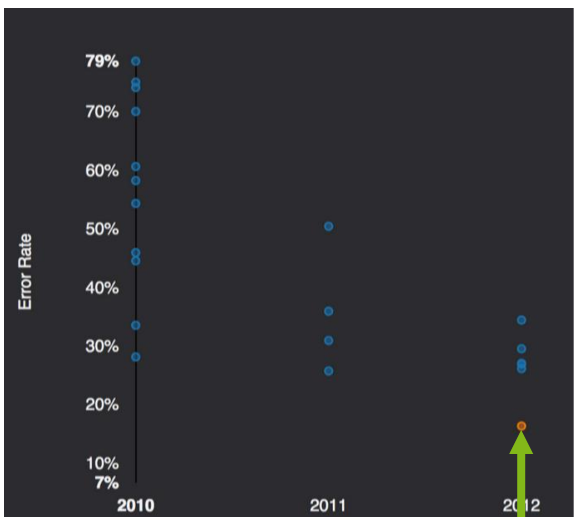

A. Krizhevsky verwendet als erster ein sog. «Deep Neural Network» (CNN)

Was ist passiert?
Der ImageNet Wettbewerb

1000 Kategorien
1 Mio. Beispiele

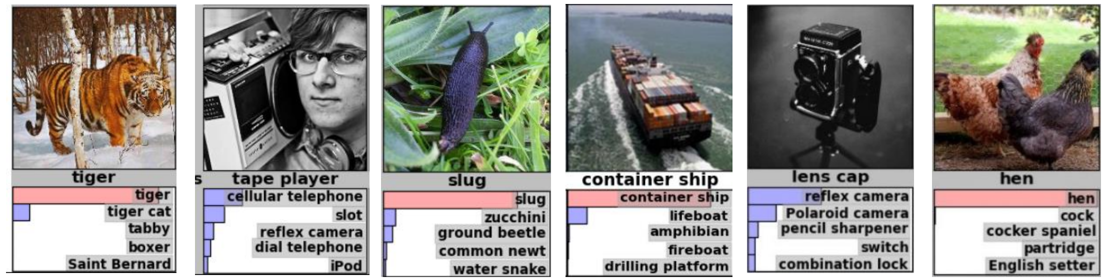

…

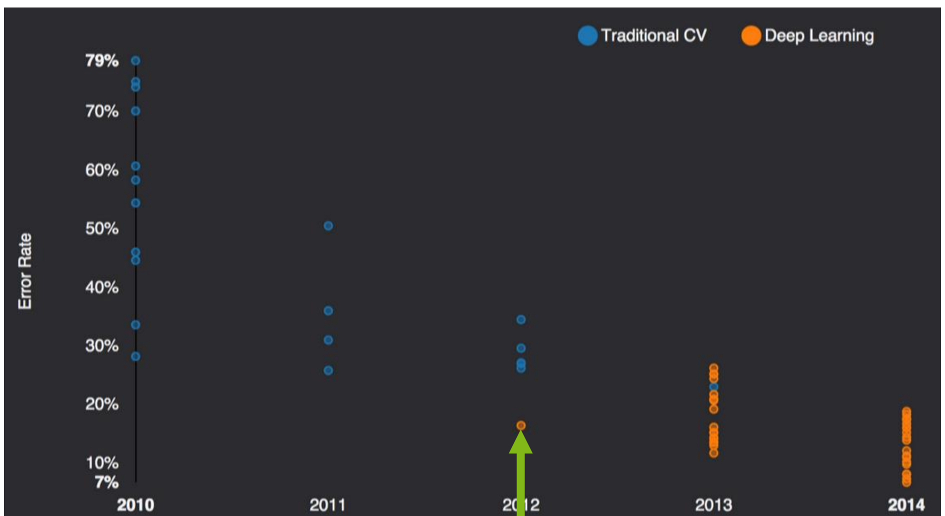

A. Krizhevsky verwendet als erster ein sog. «Deep Neural Network» (CNN)

Was ist passiert?
Der ImageNet Wettbewerb

1000 Kategorien
1 Mio. Beispiele

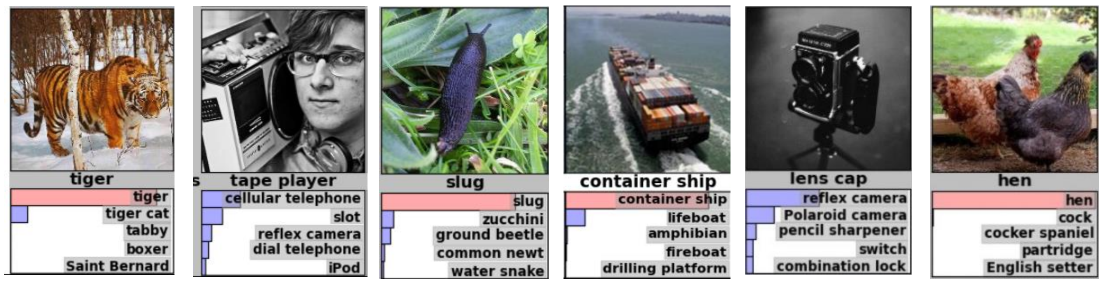

…

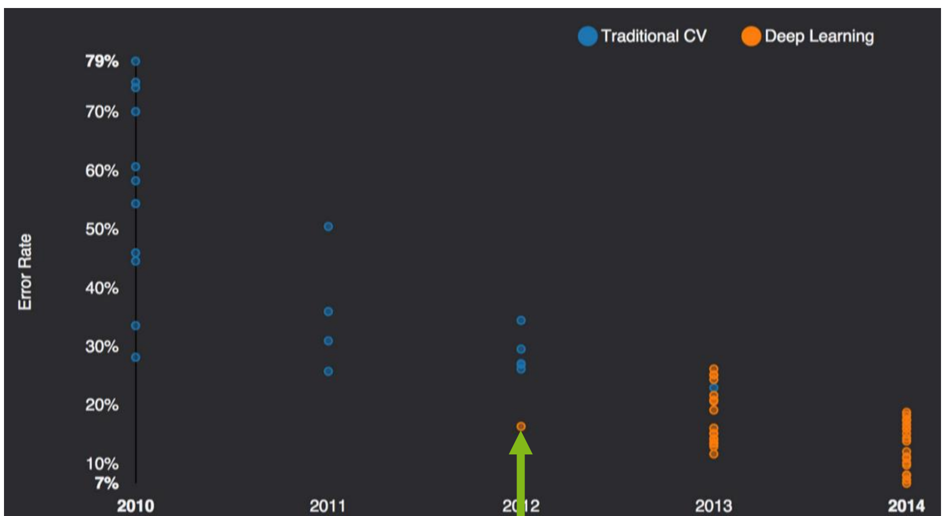

A. Krizhevsky verwendet als erster ein sog. «Deep Neural Network» (CNN)

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)

Was?→ Wie? → Wohin?

2

Wie geht das?
(Was denken denkende Maschinen?)

Grundlage

Induktives überwachtes Lernen

Annahme

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

Methode

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

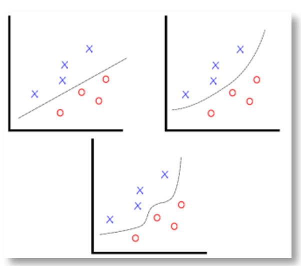

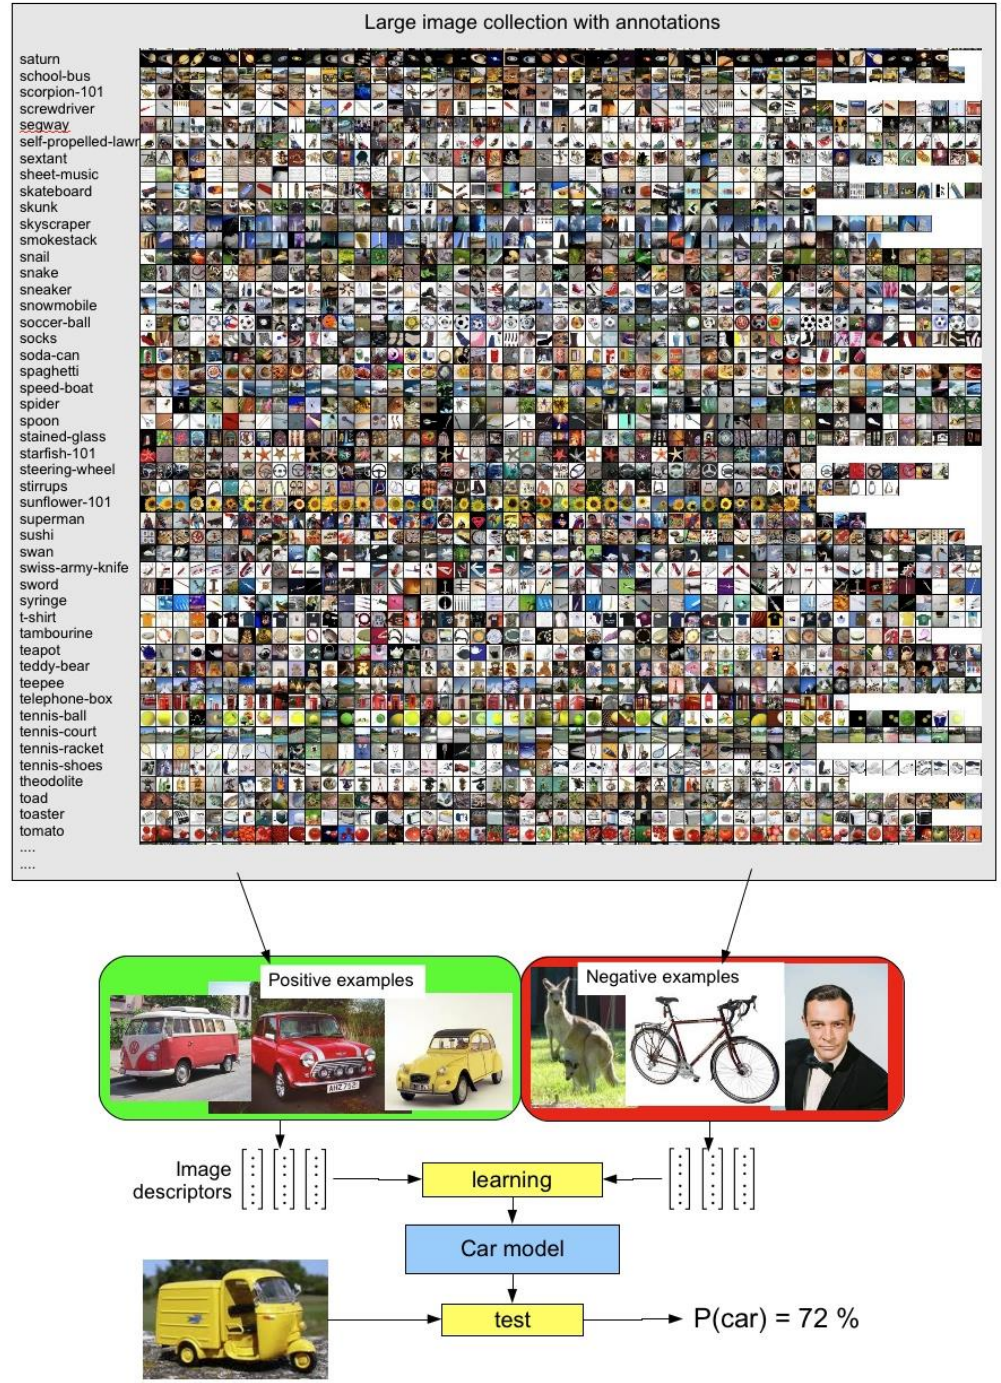

Quelle: http://lear.inrialpes.fr/job/postdoc-large-scale-classif-11-img/attribs_patchwork.jpg

Suche der Parameter einer Funktion??

Neuron

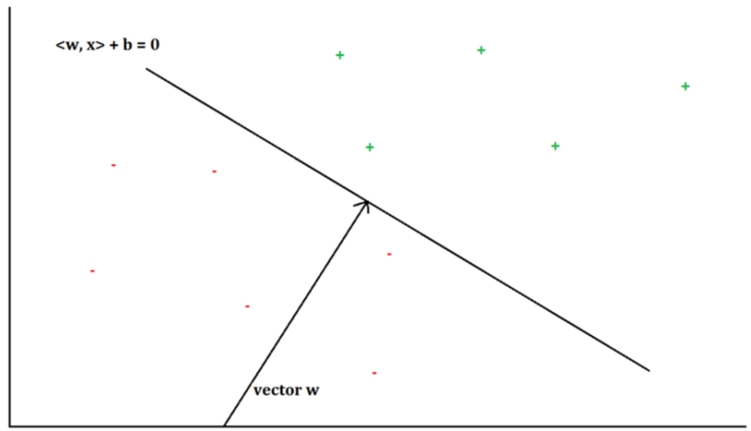

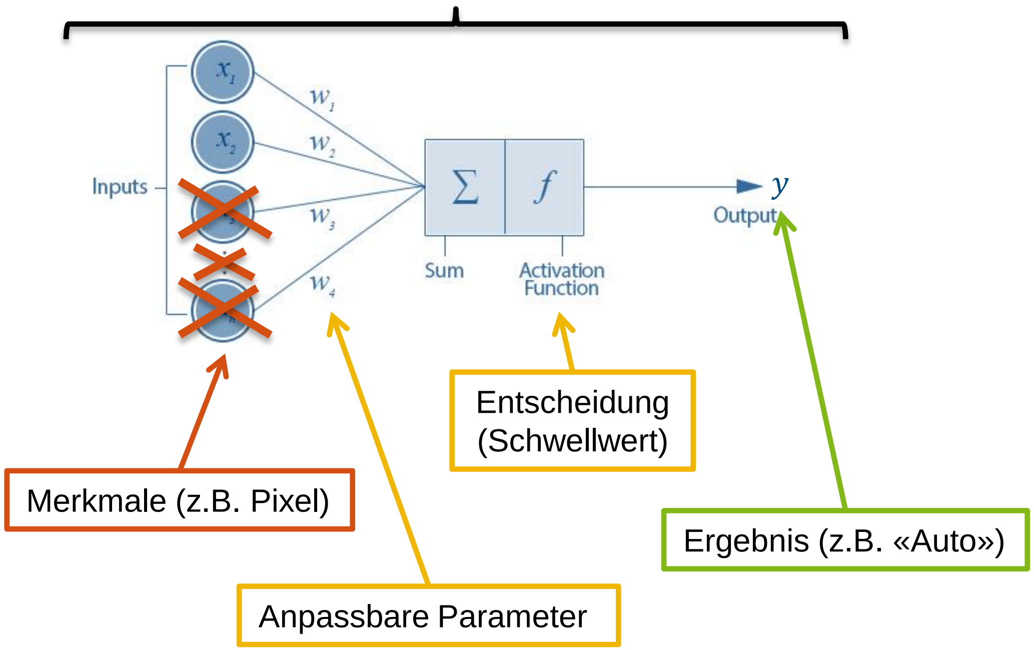

Suche der Parameter einer Funktion??

Neuron

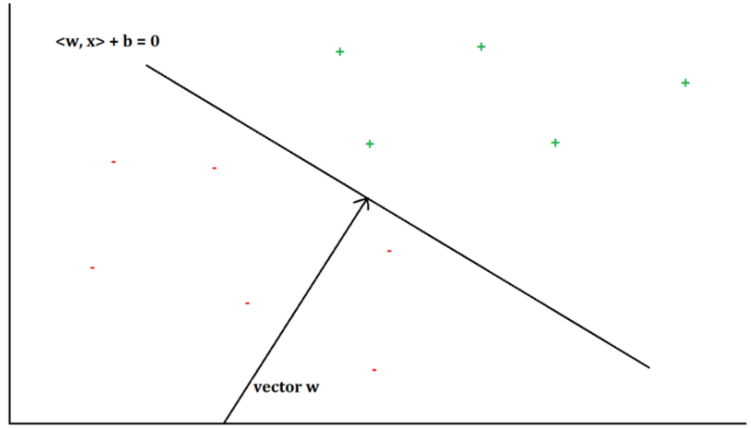

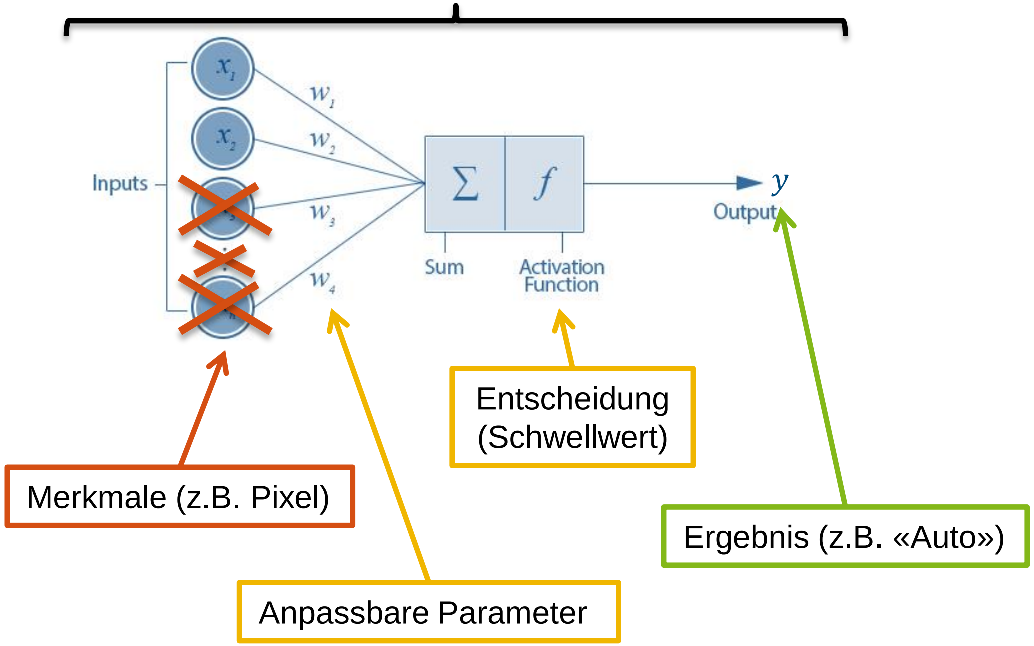

Neuronales Netz

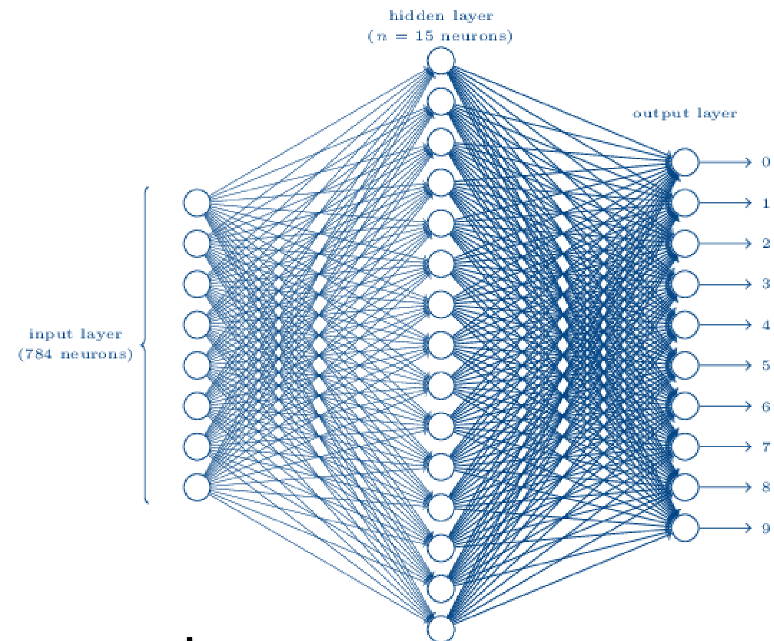

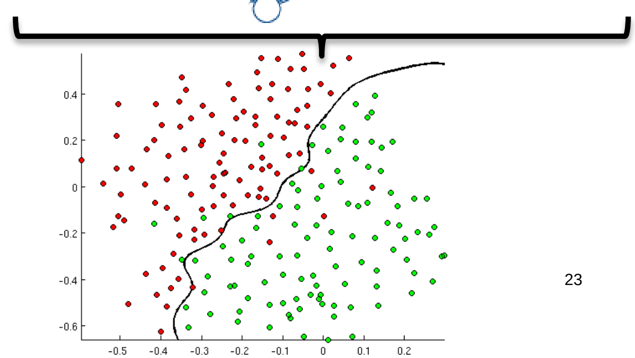

Idee: Mehr Tiefe zum Lernen von Merkmalen

Klassische Bild-verarbeitung

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

(0.2, 0.4, …)

(0.4, 0.3, …)

Klassifikation
(SVM, Neuronales Netz, etc.)

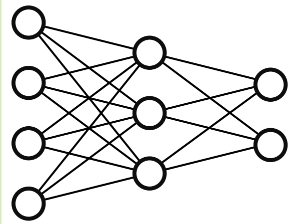

Containerschiff

Tiger

…

Mit Convolutional Neural Networks (CNNs)

Nimmt rohe Pixel entgegen,
Merkmale werden mitgelernt!

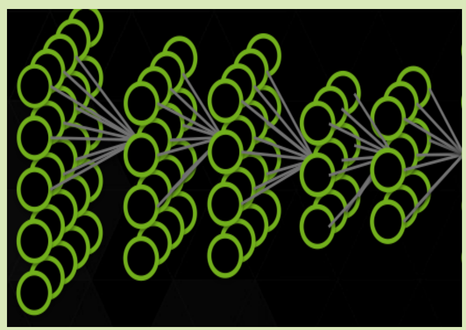

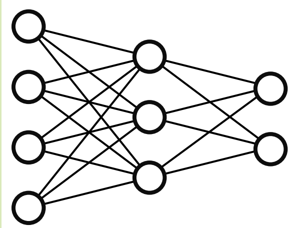

Containerschiff

Tiger

…

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

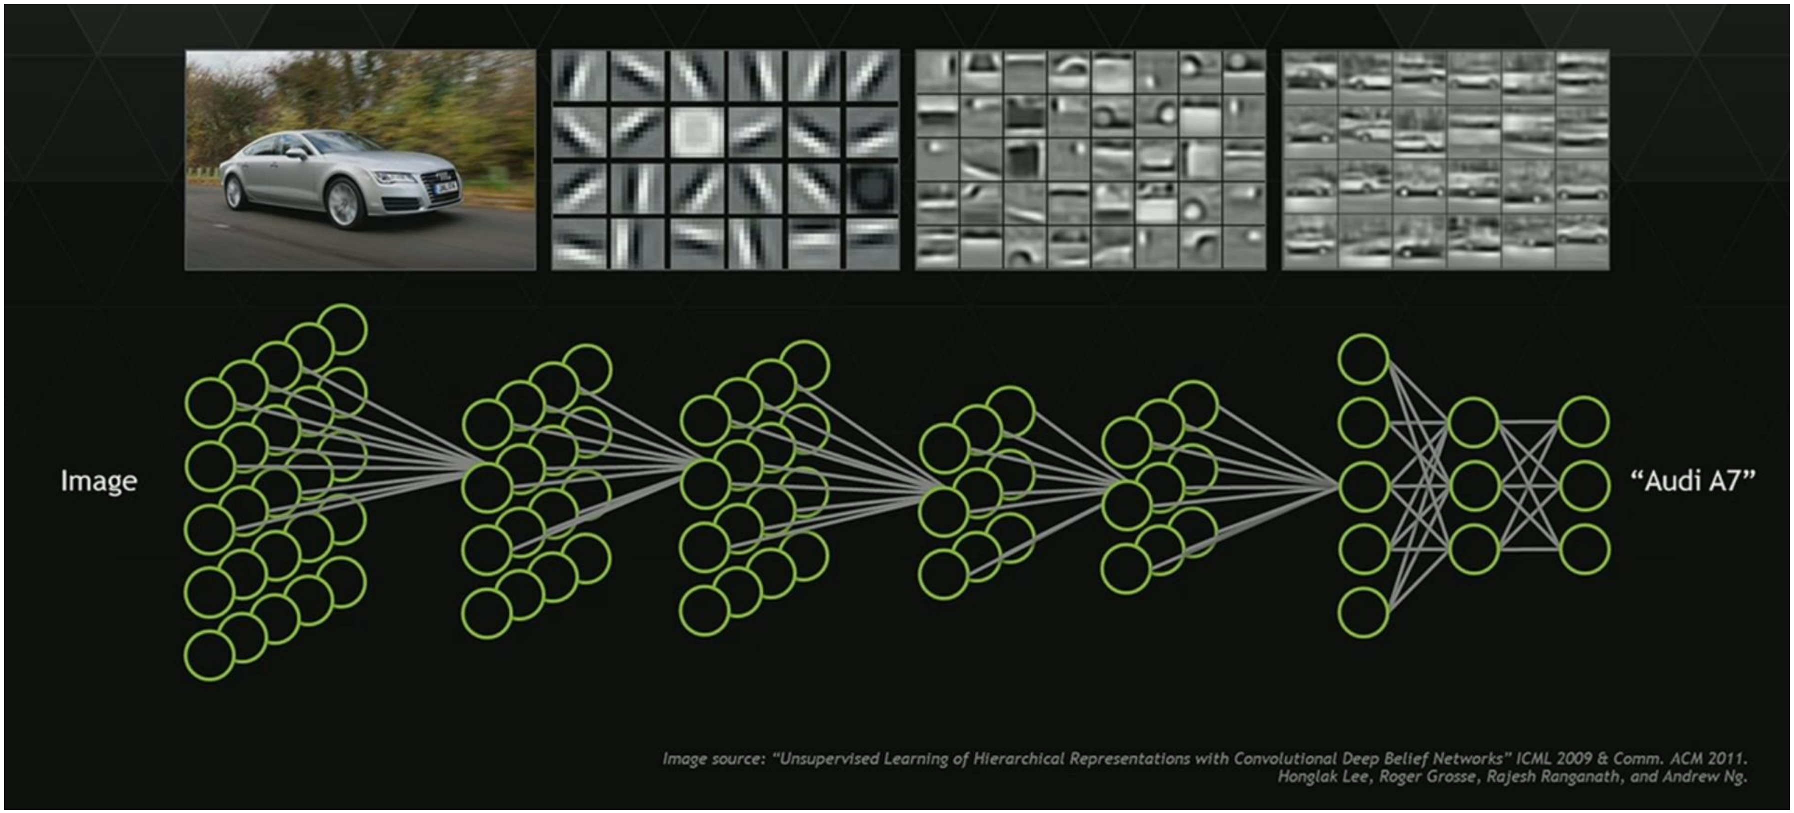

Quelle: https://www.pinterest.com/explore/artificial-neural-network/

Was?→ Wie? → Wohin?

3

Wohin führt das?
(Ein Ausblick)

Was ich <u>nicht</u> erwarte

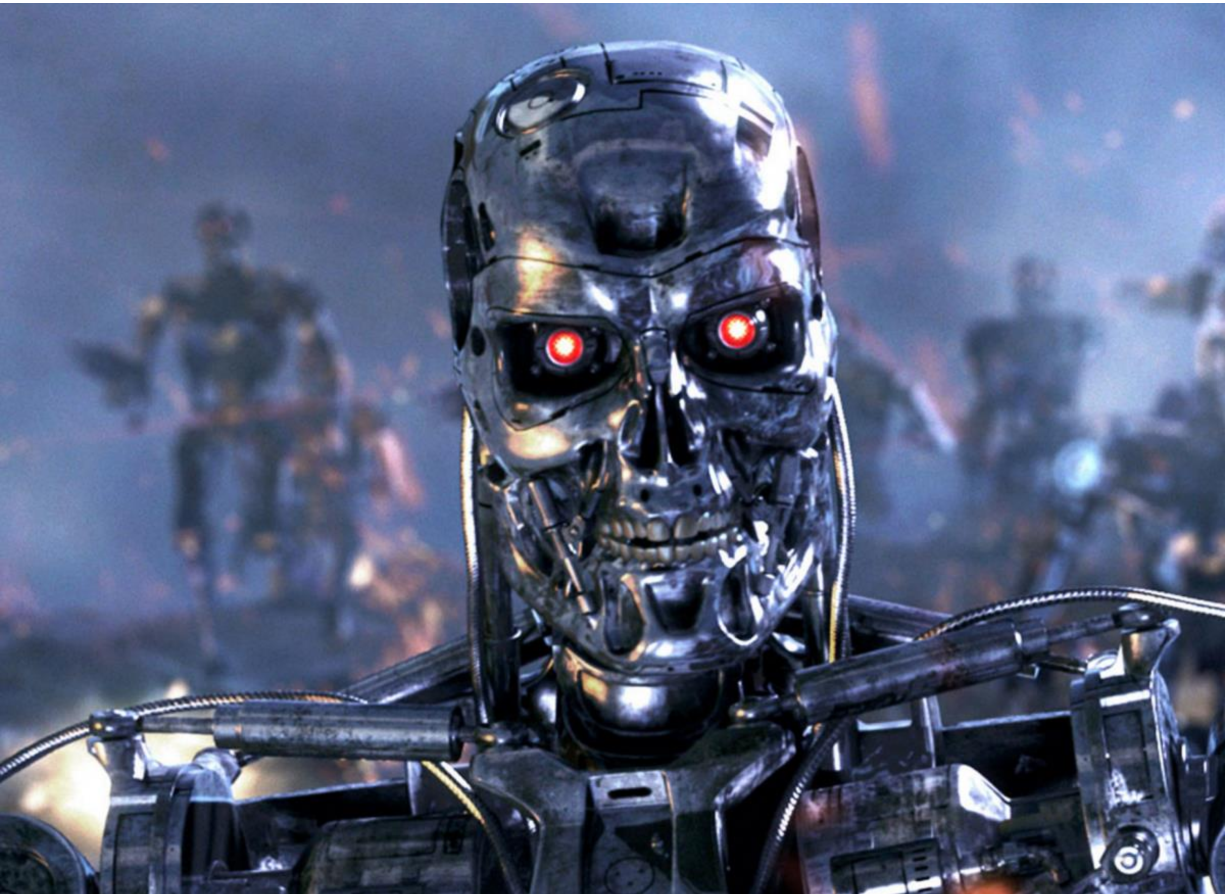

Was ich nicht erwarte

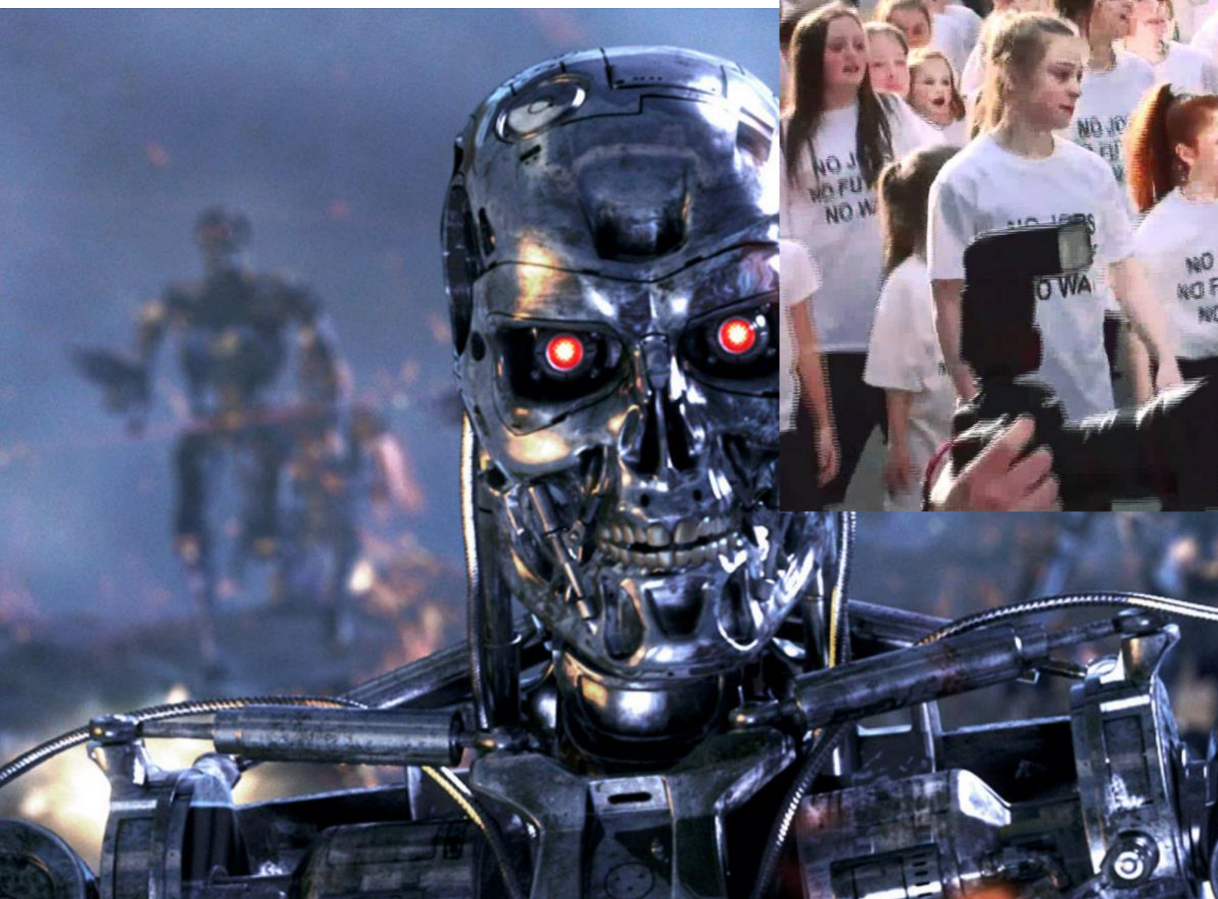

Was ich <u>nicht</u> erwarte

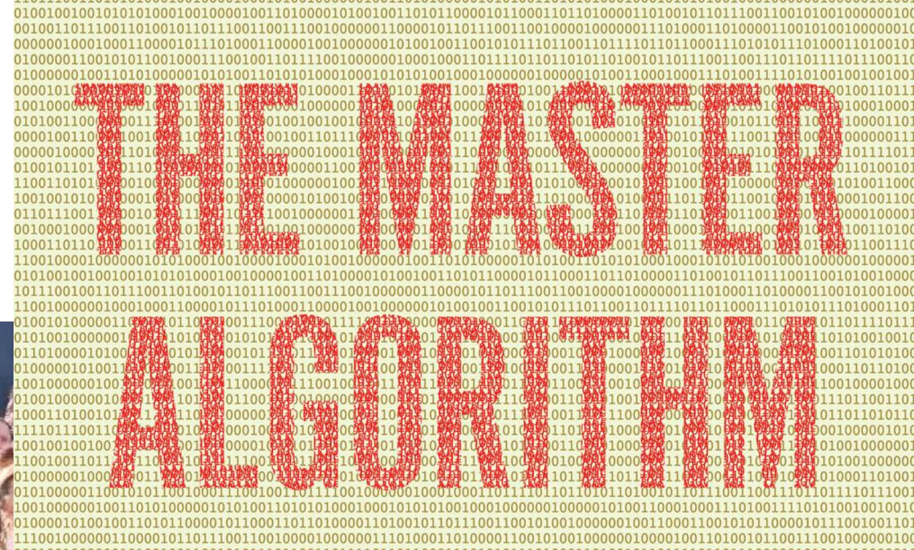

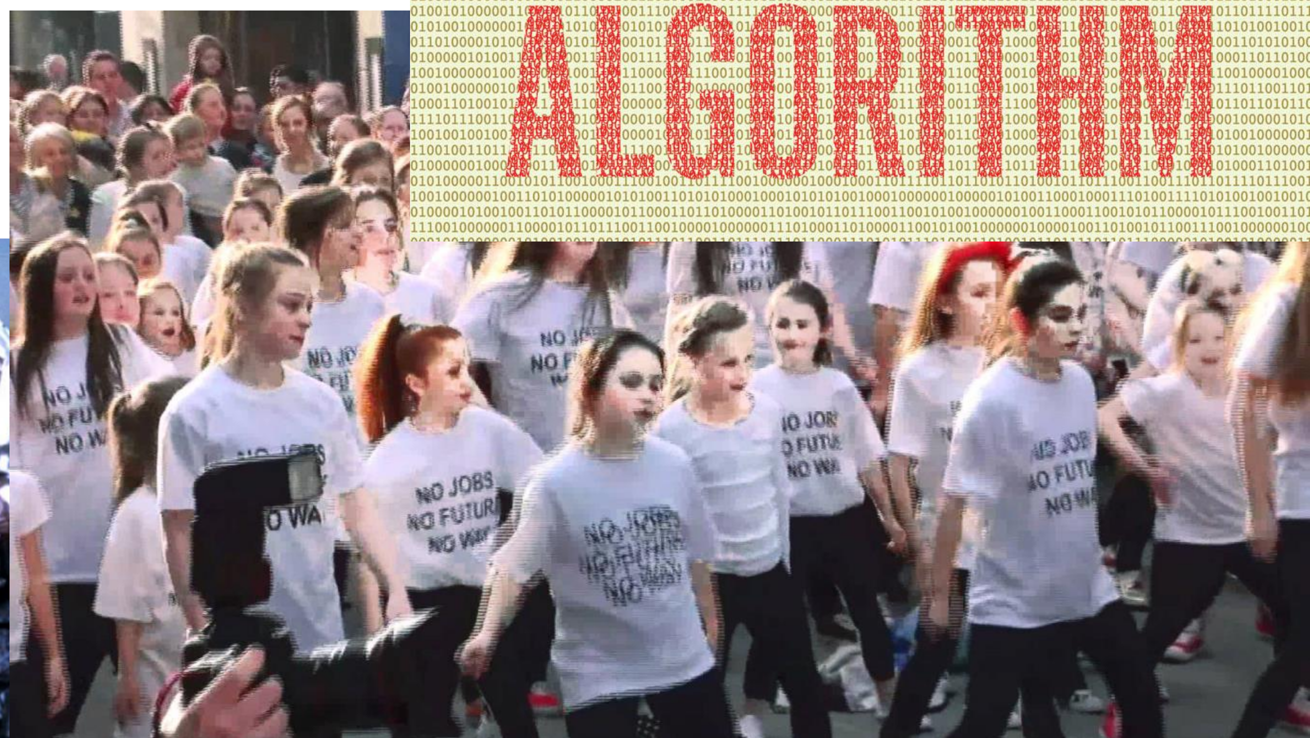

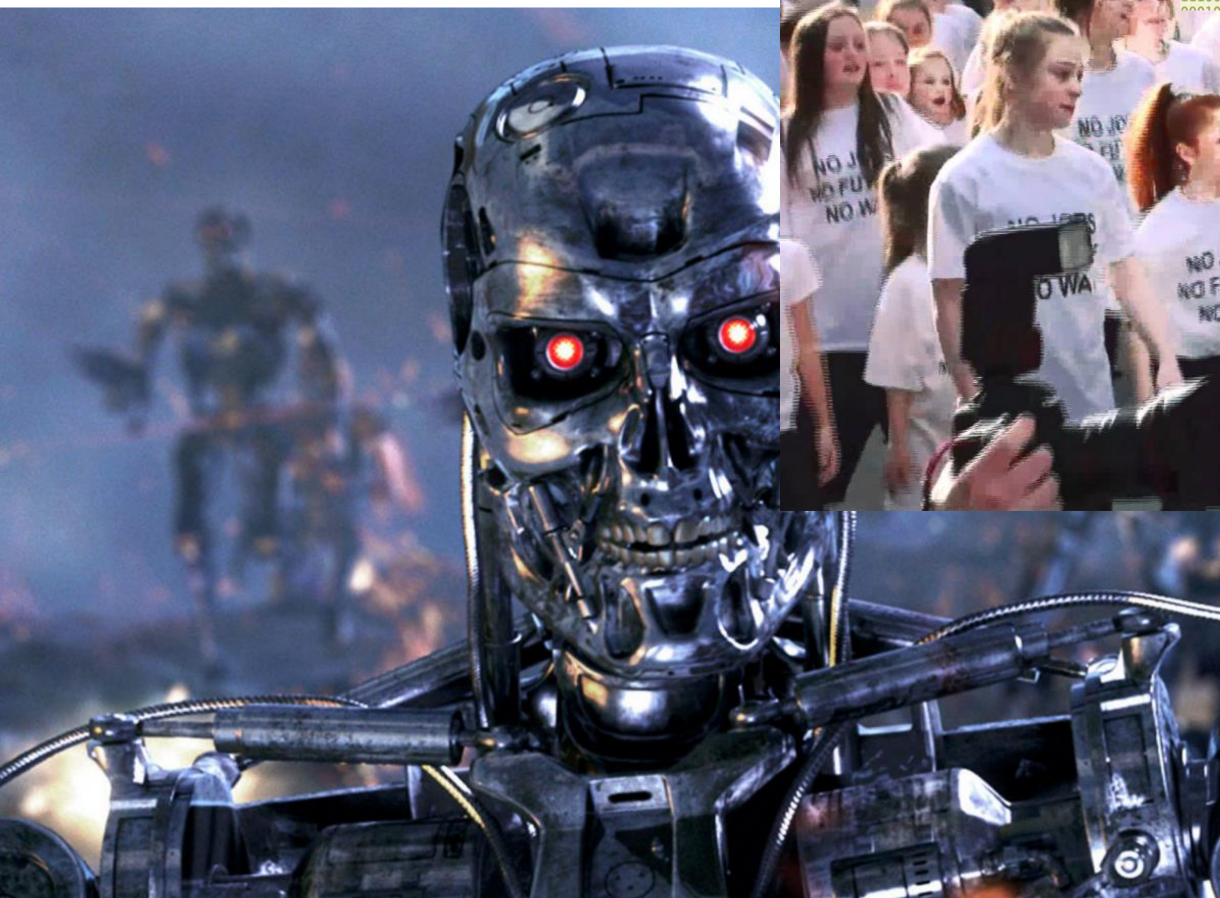

Was ich erwarte

MODERN DATA SCIENTIST

Data Scientist, the sexiest job of the 21th century, requires a mixture of multidisciplinary skills ranging from an intersection of mathematics, statistics, computer science, communication and business. Finding a data scientist is hard. Finding people who understand who a data scientist is, is equally hard. So here is a little cheat sheet on who the modern data scientist really is.

MATH & STATISTICS

- ☆ Machine learning
- ☆ Statistical modeling
- ☆ Experiment design
- ☆ Bayesian inference
- ☆ Supervised learning: decision trees, random forests, logistic regression
- ☆ Unsupervised learning: clustering, dimensionality reduction
- ☆ Optimization: gradient descent and variants

PROGRAMMING & DATABASE

- ☆ Computer science fundamentals
- ☆ Scripting language e.g. Python
- ☆ Statistical computing packages, e.g., R
- ☆ Databases: SQL and NoSQL
- ☆ Relational algebra
- ☆ Parallel databases and parallel query processing
- ☆ MapReduce concepts
- ☆ Hadoop and Hive/Pig
- ☆ Custom reducers
- ☆ Experience with xaaS like AWS

DOMAIN KNOWLEDGE & SOFT SKILLS

- ☆ Passionate about the business
- ☆ Curious about data
- ☆ Influence without authority
- ☆ Hacker mindset
- ☆ Problem solver
- ☆ Strategic, proactive, creative, innovative and collaborative

COMMUNICATION & VISUALIZATION

- ☆ Able to engage with senior management
- ☆ Story telling skills
- ☆ Translate data-driven insights into decisions and actions
- ☆ Visual art design
- ☆ R packages like ggplot or lattice
- ☆ Knowledge of any of visualization tools e.g. Flare, D3.js, Tableau

Zürcher Hochschule
für Angewandte Wissenschaften

Was ich erwarte

Zürcher Hochschule
für Angewandte Wissenschaften

Was ich erwarte

Was ich erwarte

Zürcher Hochschule
für Angewandte Wissenschaften

zh
aw

Was ich erwarte

Schlussfolgerungen

- «~~Denkende~~ rechnende» Maschinen sind immer noch *inselbegabt*
- Aber: «Deep Learning» hat zu Quantensprung in *Mustererkennungsaufgaben* geführt
- Für andere Datenanalyseaufgaben sind *andere Verfahren besser* geeignet

- Angst ist unangebracht – aber Herausforderungen wollen gestaltet werden: technisch, ethisch, wirtschaftlich, gesellschaftlich

datalab
www.zhaw.ch/datalab

swiss group for artificial intelligence and cognitive science
SGAICO

data service
Swiss Alliance for Data-Intensive Services

Mehr zu mir:

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

Mehr zum Thema:

- KI: http://www.s-i.ch/en/fachgruppen-und-sektionen/sgaico/
- Verband Data & Service Science: www.data-service-alliance.ch
- Gemeinsame Projekte: datalab@zhaw.ch

➔ Fragen Sie gerne an.

ANHANG

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

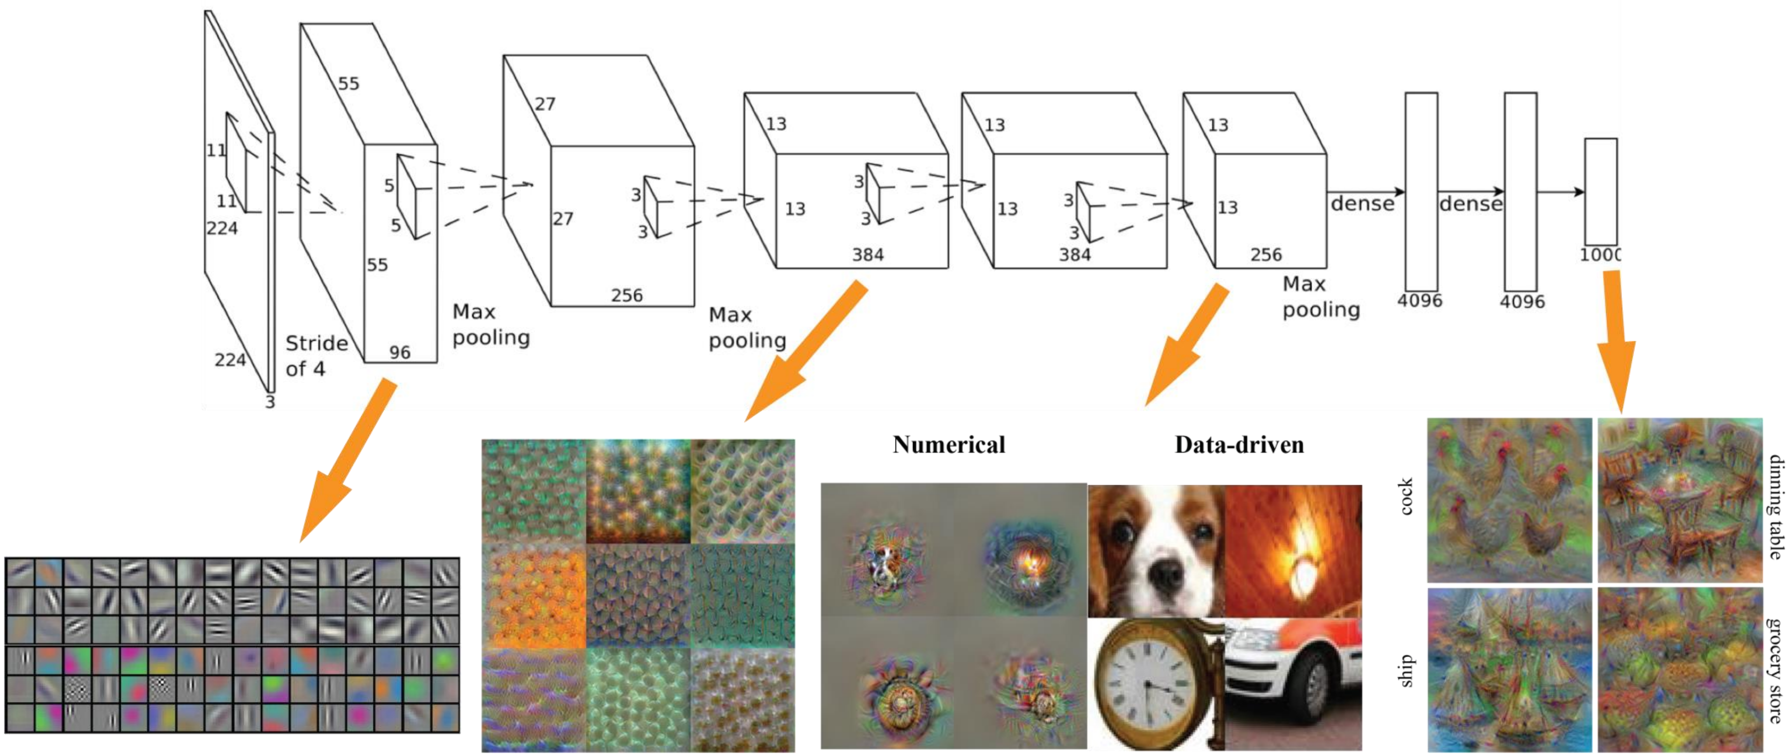

Conv 1: Edge+Blob

Conv 3: Texture

Conv 5: Object Parts

Fc8: Object Classes

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