Zürcher Hochschule
für Angewandte Wissenschaften

Deep Learning from an IT perspective

Industry session @ 6th Richmond IT Forum, Bad Ragaz, Sep 28, 2020

Thilo Stadelmann

What is AI and DL?
Examples for successful DL deployments
Lessons learned from an IT perspective

datalab
www.zhaw.ch/datalab

Image source: https://www.resortragaz.ch/de/aktivitaeten-und-events/erlebnisse

What→ Examples?→ Lessons learned

1

What is AI and Deep Learning?

What is AI?

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

What belongs to AI?

An incomplete view of its subdisciplines

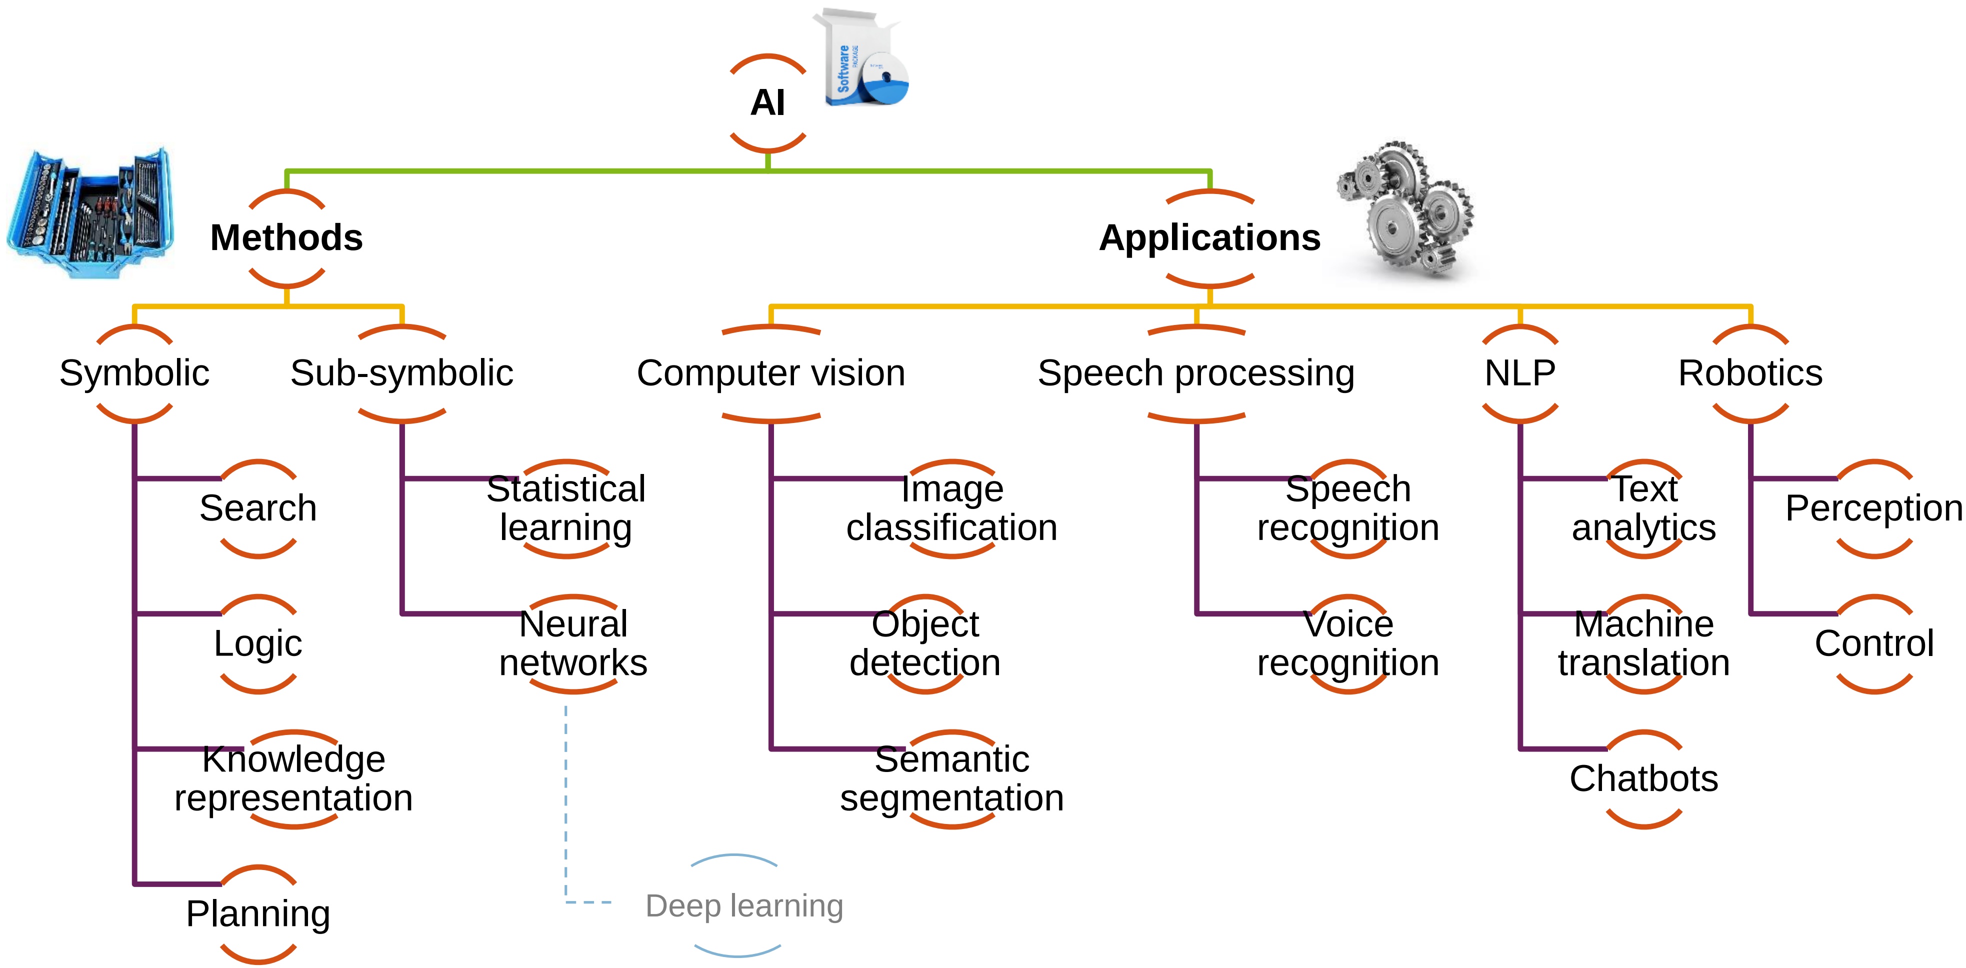

Zürcher Hochschule
für Angewandte Wissenschaften

Why is AI big *now*?

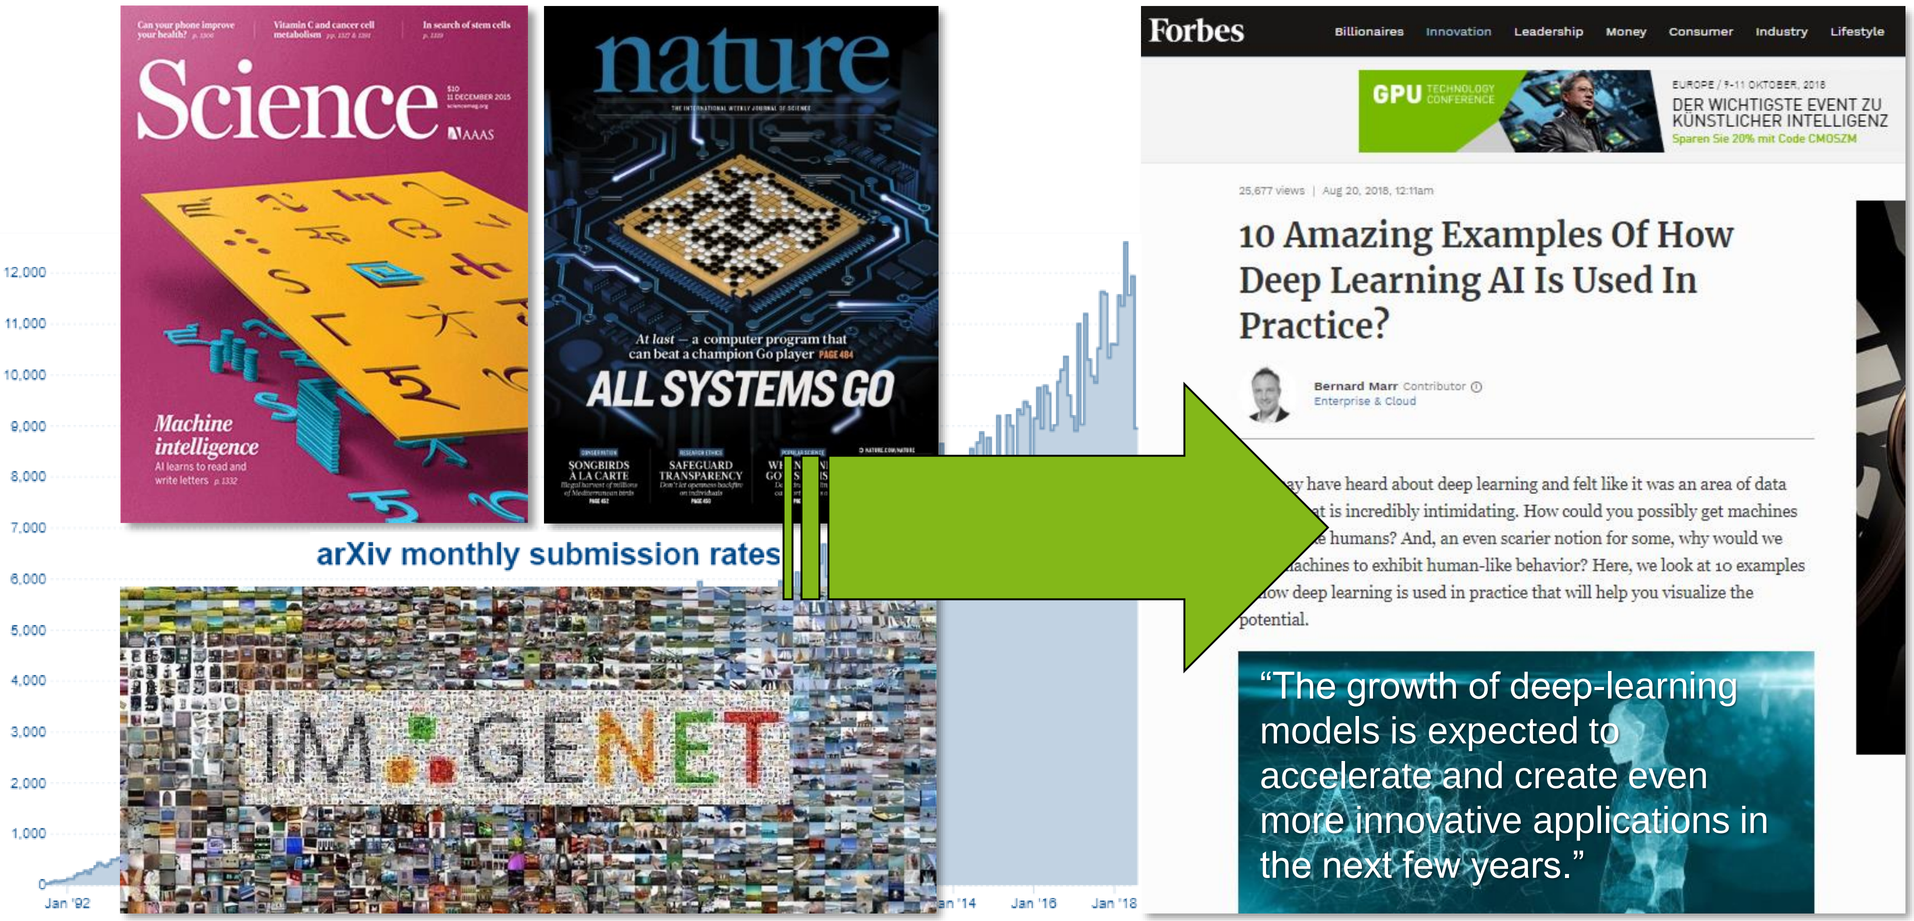

What’s the big deal about deep learning?
Adding depth to learn features automatically

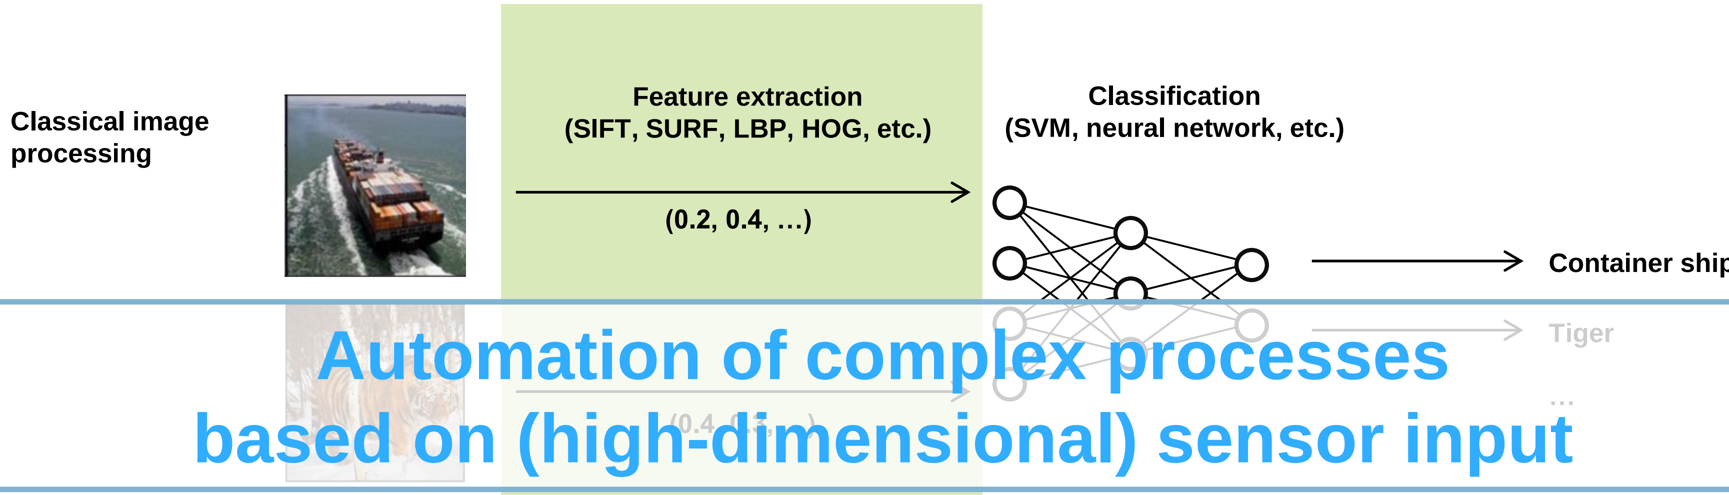

Automation of complex processes based on (high-dimensional) sensor input

What→ Examples?→ Lessons learned

2

Examples for successful DL deployments

Zürcher Hochschule
für Angewandte Wissenschaften

Examples of «AI» in the media in recent years

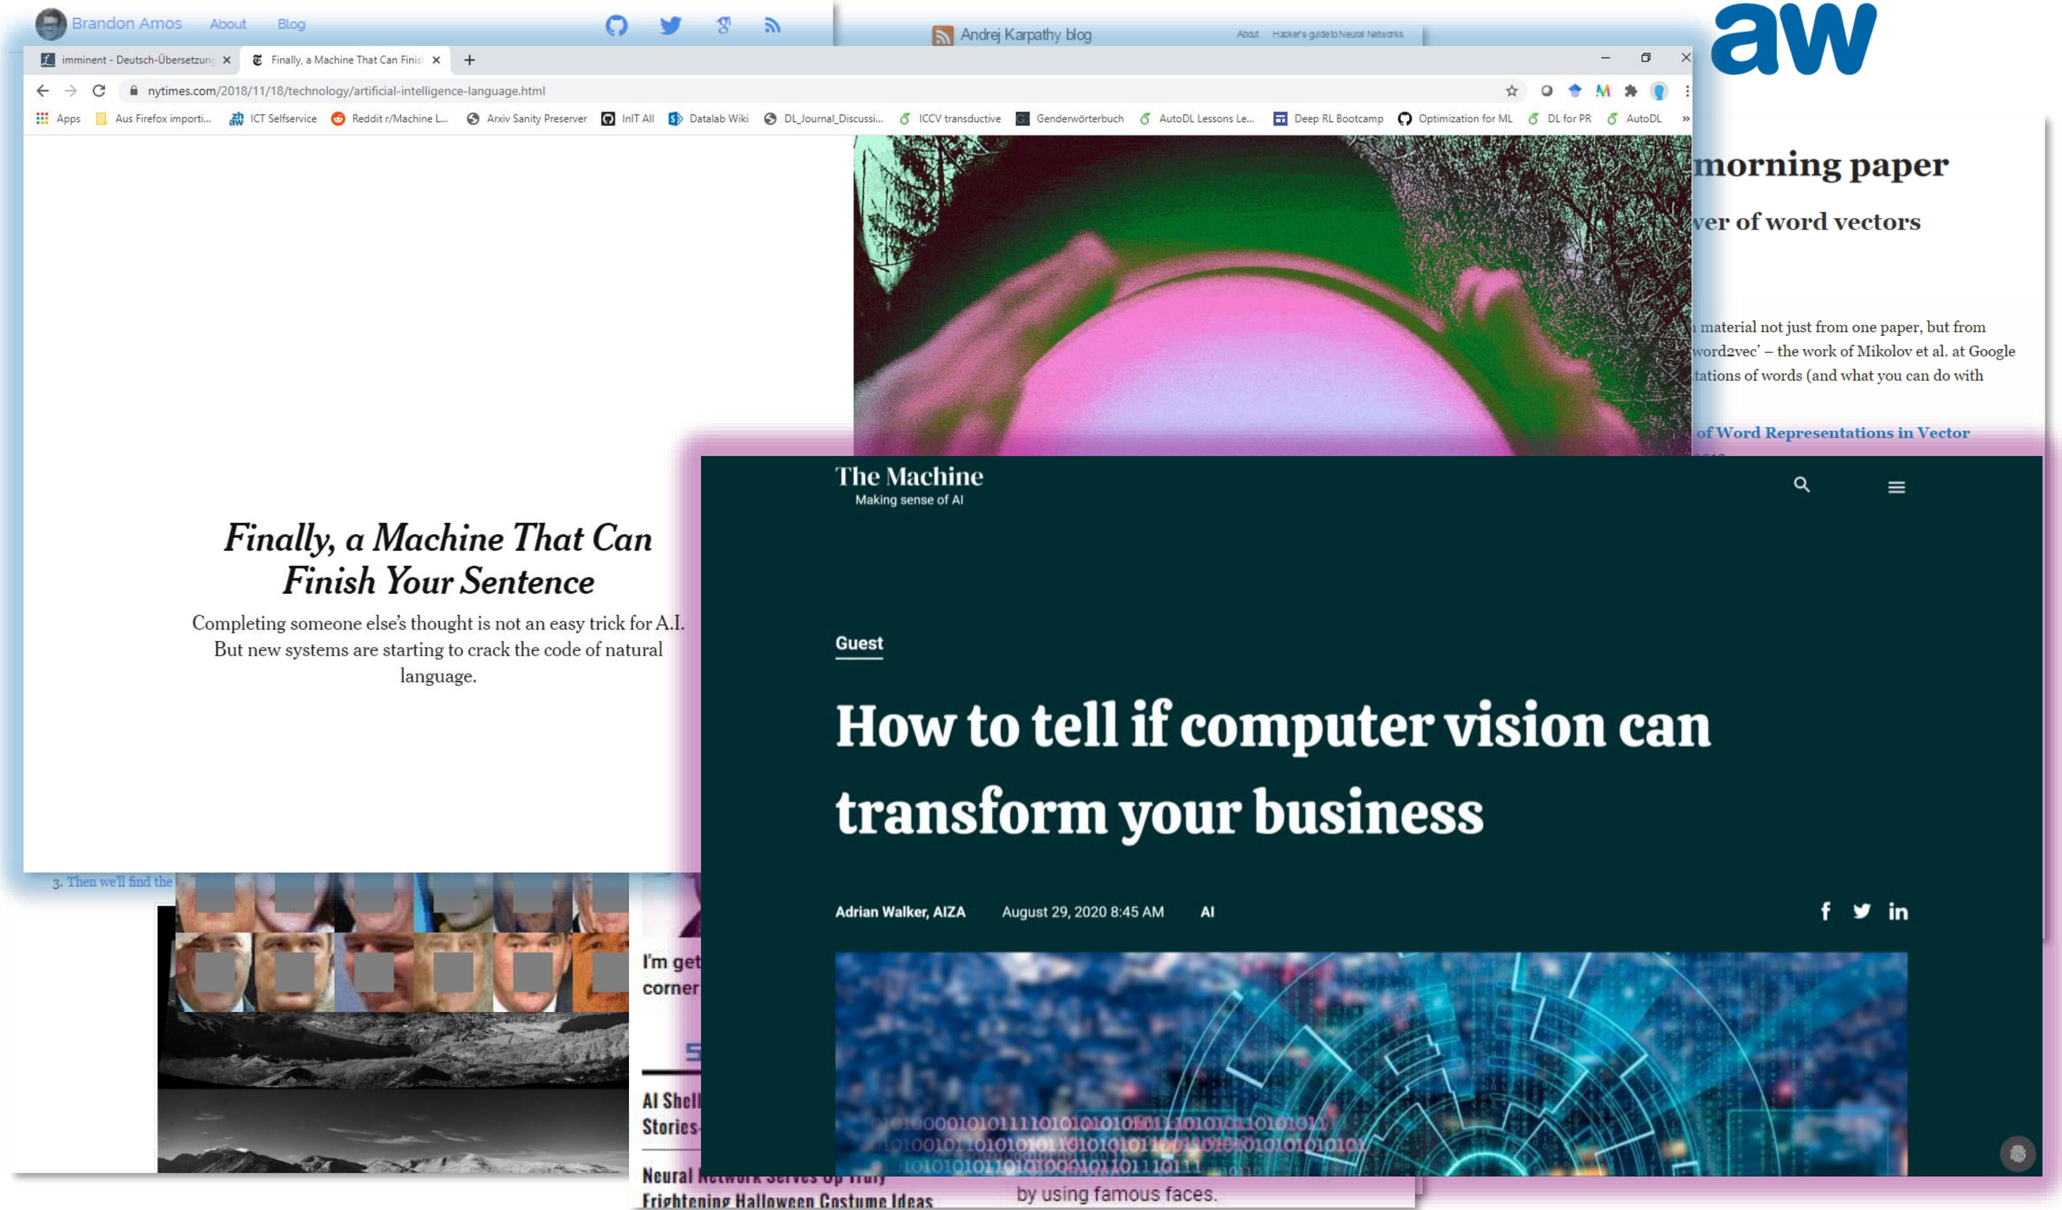

Use case 1: print media monitoring

ARGUS DATA INSIGHTS
WISSEN ZUM ERFOLG

Schweizerische Eidgenossenschaft
Confédération suisse
Confederazione Svizzera
Confederaziun svizra

Swiss Confederation

Innosuisse – Swiss Innovation Agency

Task


International.

Nachrichten

Spionage für den Erzfeind Iran

Israelischer Ex-Minister arbeitete als Agent für die Mullahs. Jetzt droht ihm lebenslänglich

Vermögen beschlagnahmt

Asylbewerber können bleiben

Klage von Le Pen abgelehnt

Nordkoreanischer Diktator zu Besuch in Peking

Challenge


Sport Blick 15

«Steven hat sich alles selber beigebracht»

Unsere Nati ist gegen Serbien unter Druck

Verlieren verboten!

Transfer TICKER

Liverpool will Yann Sommer!

Nuisance


Das Tages-Horoskop

Liebling der Sterne

15,1 Millionen

Sind Sie der nächste Lotto-König?

Jungfrau

Waage

Skorpion

Schütze

Steinbock

Wassermann

Fische

Widder

Stier

Zwillinge

Krebs

SCHWEDEN-RÄTSEL

SUDOKU

Wochenpreis: 1 x sieben Nächte für 2 Personen, inkl. HP, im ****Seehotel Pilatus Hergiswil im Wert von 3000 Franken!

Print media monitoring – ML solution

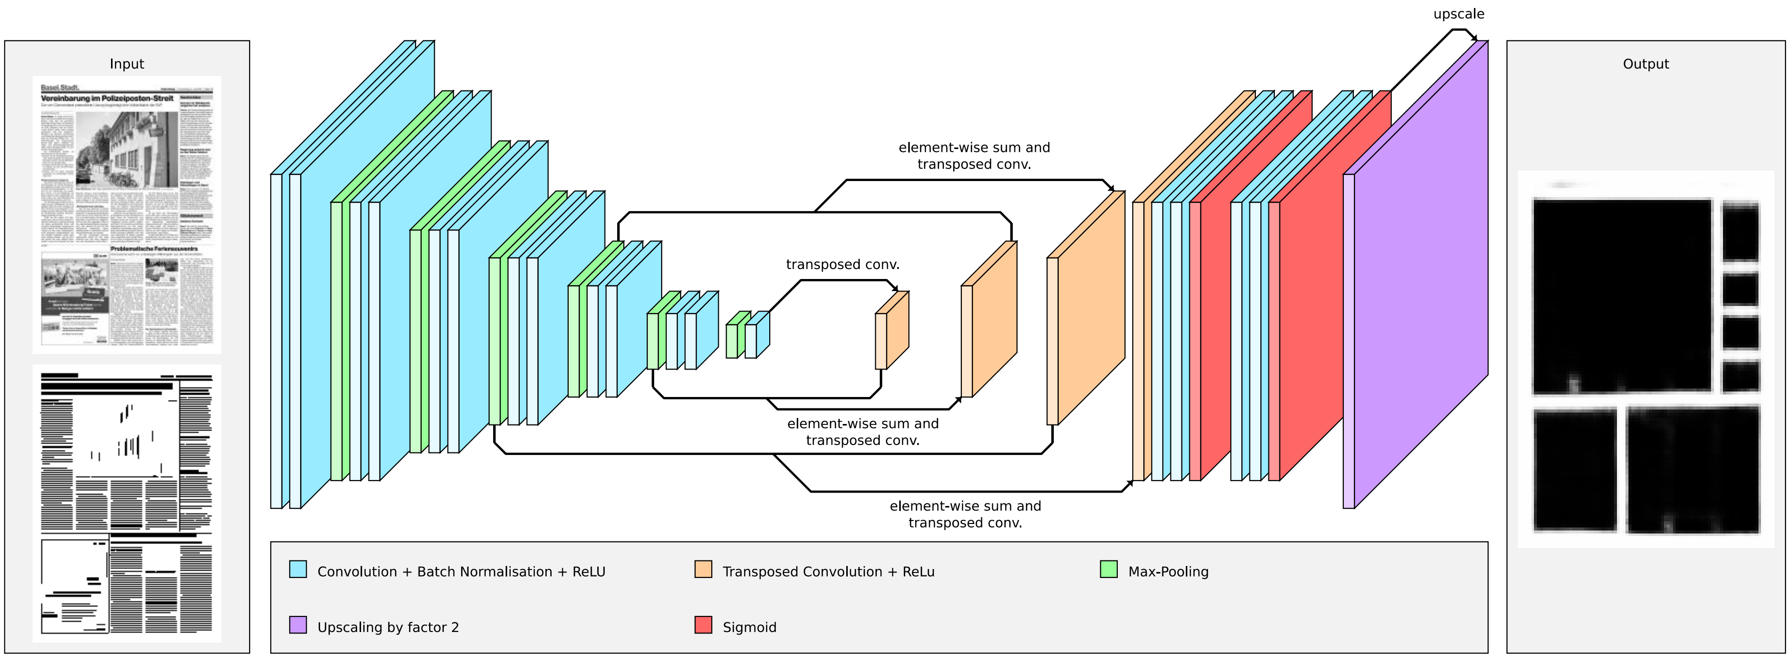

Meier, Stadelmann, Stampfli, Arnold & Cieliebak (2017). *«Fully Convolutional Neural Networks for Newspaper Article Segmentation»*. ICDAR'2017.
Stadelmann, Tolkachev, Sick, Stampfli & Dürr (2018). *«Beyond ImageNet - Deep Learning in Industrial Practice»*. In: Braschler et al., «Applied Data Science»,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

Print media monitoring – deployment

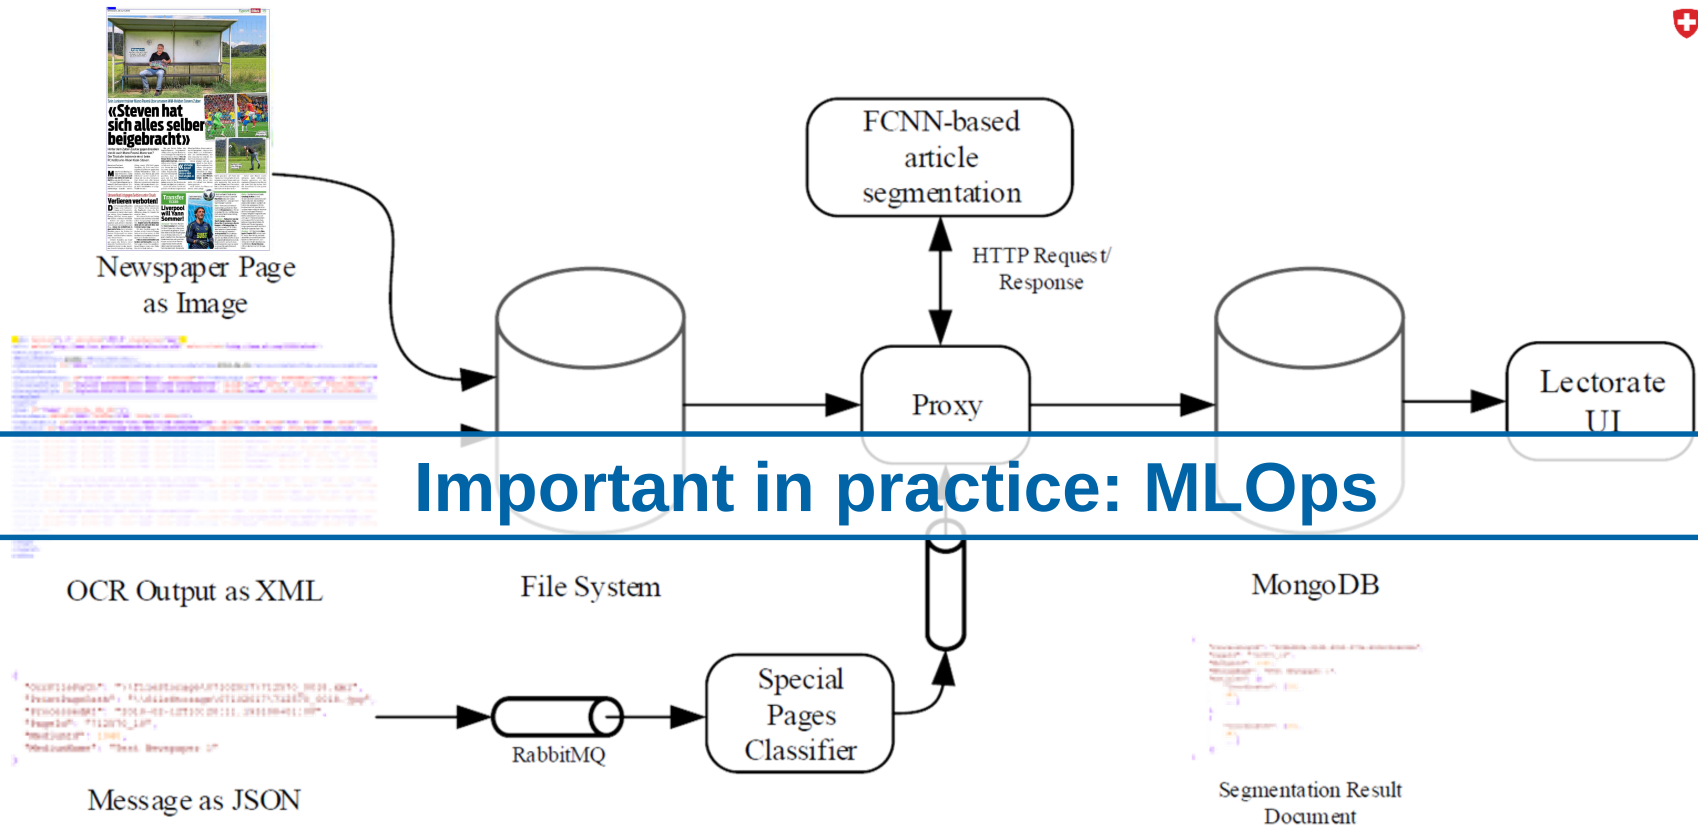

Important in practice: MLOps

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

Use case 2: symbol detection

Zürcher Hochschule
für Angewandte Wissenschaften

SCOREPAD

Schweizerische Eidgenossenschaft
Confédération suisse
Confederazione Svizzera
Confederaziun svizra

Swiss Confederation

Innosuisse – Swiss Innovation Agency

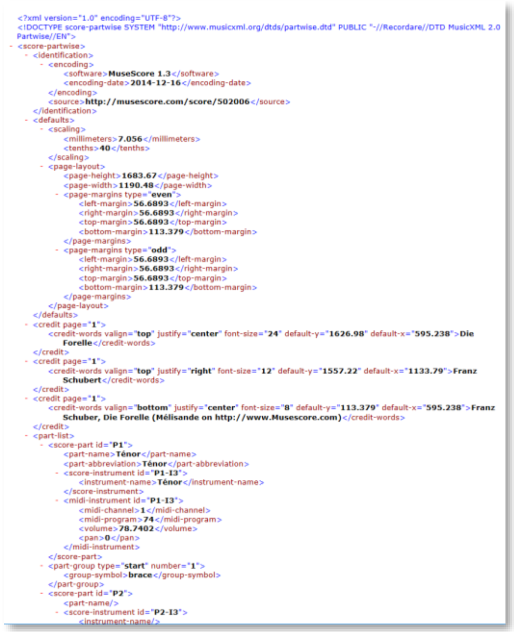

```xml
<?xml version="1.0" encoding="UTF-8"?>
<!DOCTYPE score-partwise SYSTEM "http://www.musicxml.org/dtds/partwise.dtd" PUBLIC "-//Recordare//DTD MusicXML 2.0 Partwise//EN">
- <score-partwise>
  - <identification>
    - <encoding>
        <software>MuseScore 1.3</software>
        <encoding-date>2014-12-16</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/>
```


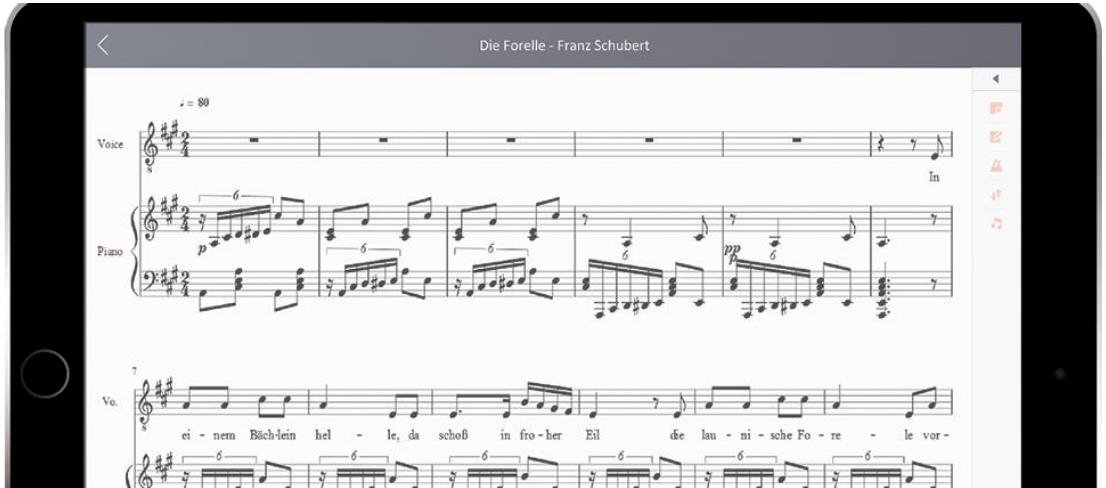

Symbol detection – challenges & solutions

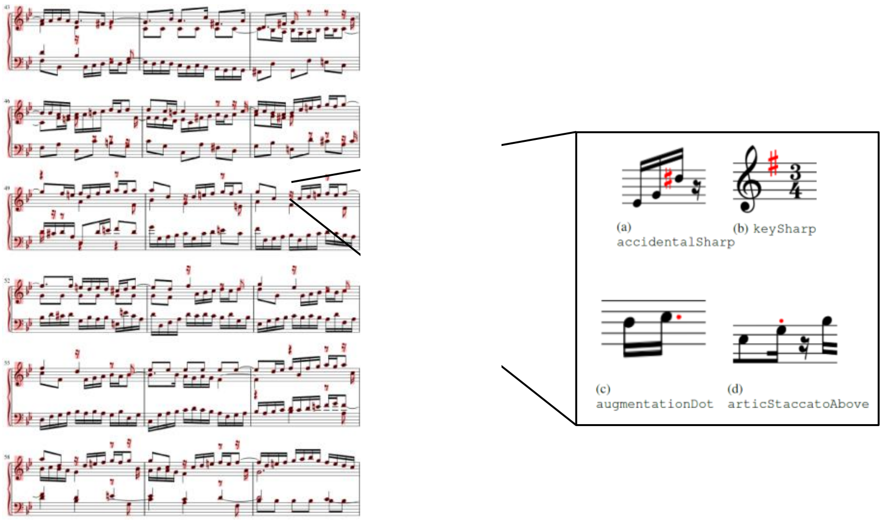

,

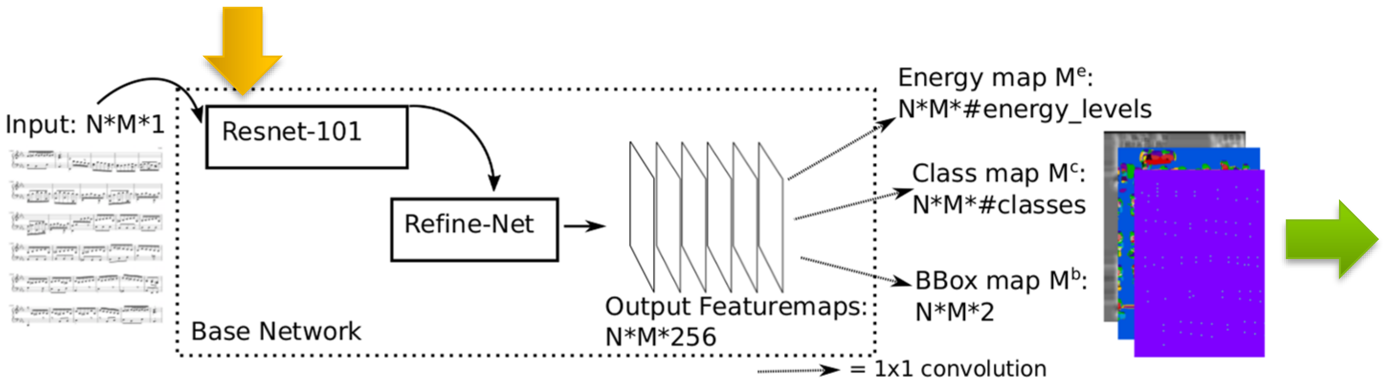

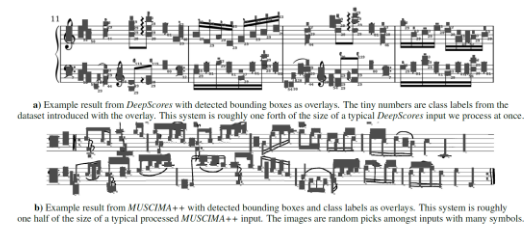

Tuggener, Elezi, Schmidhuber, Pelillo & Stadelmann (2018). *«DeepScores – A Dataset for Segmentation, Detection and Classification of Tiny Objects»*. ICPR'2018.
Tuggener, Elezi, Schmidhuber & Stadelmann (2018). *«Deep Watershed Detector for Music Object Recognition»*. ISMIR'2018.

Symbol detection – industrialization

Current results on class imbalance and robustness challenges

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

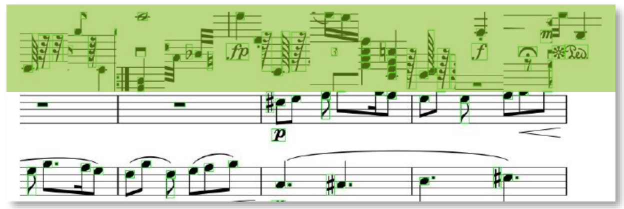

2. Put additional effort (and compute) into hyperparameter **tuning** and **longer training**
3. Trained also on scanned (more **real-worldish**) scores

Sufficient condition: lots of tuning

➔ **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: SotA 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.

What→ Examples?→ Lessons learned

3

Lessons learned from an IT perspective

Basis for disruption: automation „at scale“

Or: “digital transformation” refers to a shift in all aspects of society, driven/enabled by this small set of technologies

AI

Massively enhanced automation depth through progress in pattern recognition

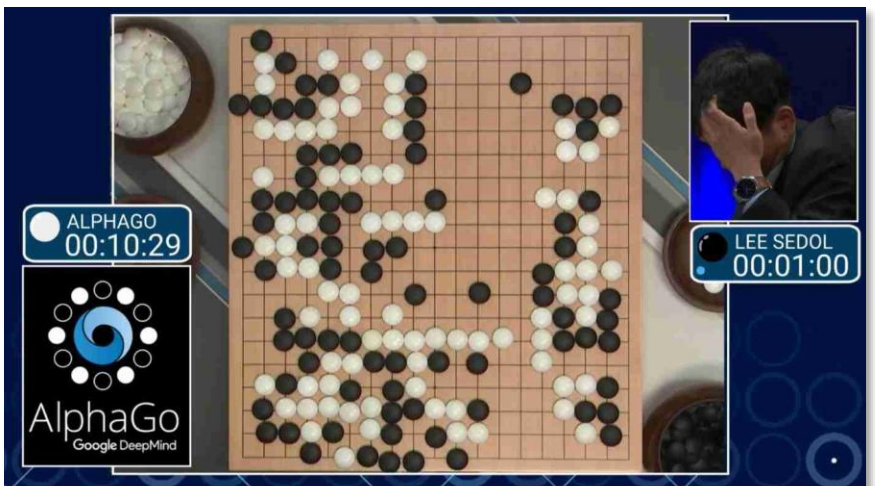

CLOUD COMPUTING

No need to invest into (IT) infrastructure anymore before entering the market

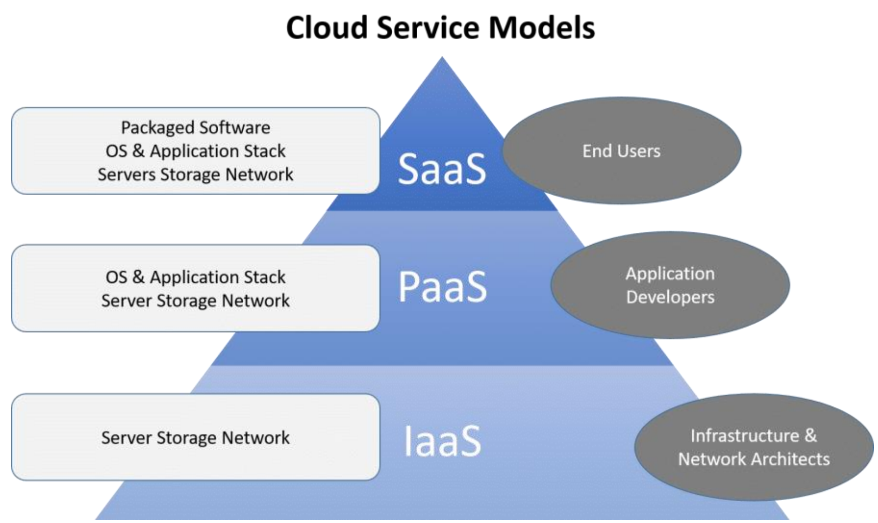

One Implication: new opportunities
…through decoupling

Zürcher Hochschule
für Angewandte Wissenschaften

size of idea ≠ size of implementing organization

…small organizations can build **whatever they want**
(given know-how, data and an interesting business case)

the technology is sector-independent

…enabling **new** alliances and co-operations

Risks through AI?

- AI per definition is a “**dual use technology**”
  → see report by Brundage et al., 2018
- But: “**natural stupidity**” is the more imminent threat
- **AI ethics** and explainable AI became mainstream and hot research topics in the recent years – not because of intolerable risks, but because of:

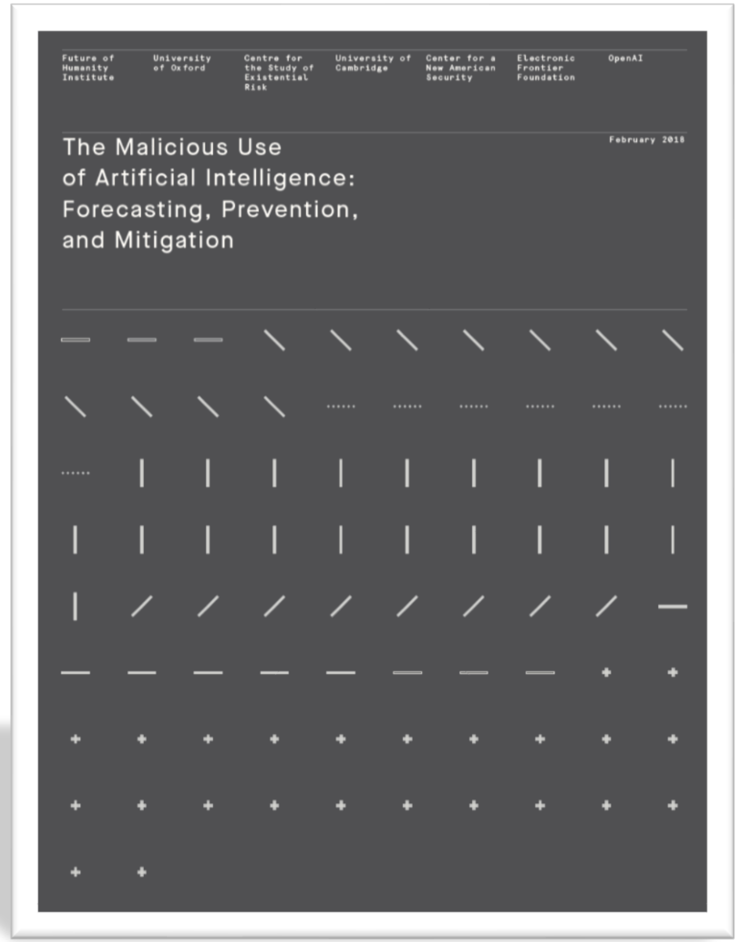

Future of Humanity Institute | University of Oxford | Centre for the Study of Existential Risk | University of Cambridge | Center for a New American Security | Electronic Frontier Foundation | OpenAI

February 2018

The Malicious Use of Artificial Intelligence: Forecasting, Prevention, and Mitigation

The risk of natural stupidity
…or the problem of customer satisfaction

SKYLIGHT   ABOUT US   SERVICES   BLOG

18 July 2019

Cylance, I Kill You!

Read about our Journey of dissecting the brain of a leading AI based Endpoint Protection Product, culminating in the creation of a universal bypass

TL;DR

AI applications in security are clear and potentially useful, however AI based products offer a new and unique attack surface. Namely, if you could truly understand how a certain model works, and the type of features it uses to reach a decision, you would have the potential to fool it consistently, creating a universal bypass.

By carefully analyzing the engine and model of Cylance's AI based antivirus product, we identify a peculiar bias towards a specific game. Combining an analysis of the feature extraction process, its heavy reliance on strings, and its strong bias for this specific game, we are capable of crafting a simple and rather amusing bypass. Namely, by appending a selected list of strings to a malicious file, we are capable of changing its score significantly, avoiding detection. This method proved successful for 100% of the top 10 Malware for May 2019, and close to 90% for a larger sample of 384 malware.

Application-dependent risks
…or the problem of feasibility and market conformity

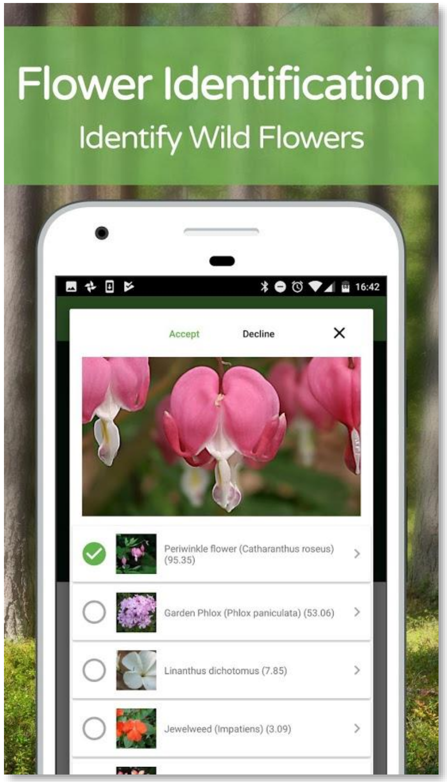

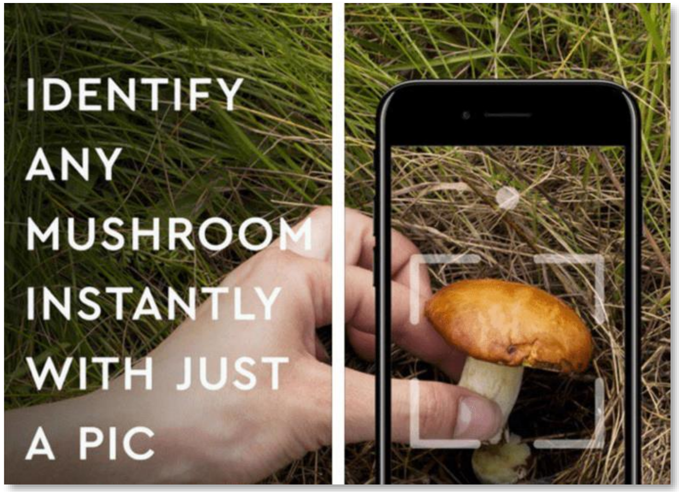

Zürcher Hochschule
für Angewandte Wissenschaften

The problem of data
Not big, but high-quality

Data is key

- Many real-world projects miss the required **quantity & quality** of data
  → even though «big data» is not needed
- **Class imbalance** needs careful dealing
  → special loss, resampling (also in unorthodox ways)
- **Unsupervised** methods need to be used creatively
- Users & label providers need to be **trained**


Source: https://www.nytimes.com/2018/11/25/business/china-artificial-intelligence-labeling.html

Zürcher Hochschule
für Angewandte Wissenschaften

The problem of compute
Training time GPT3 vs. at the edge

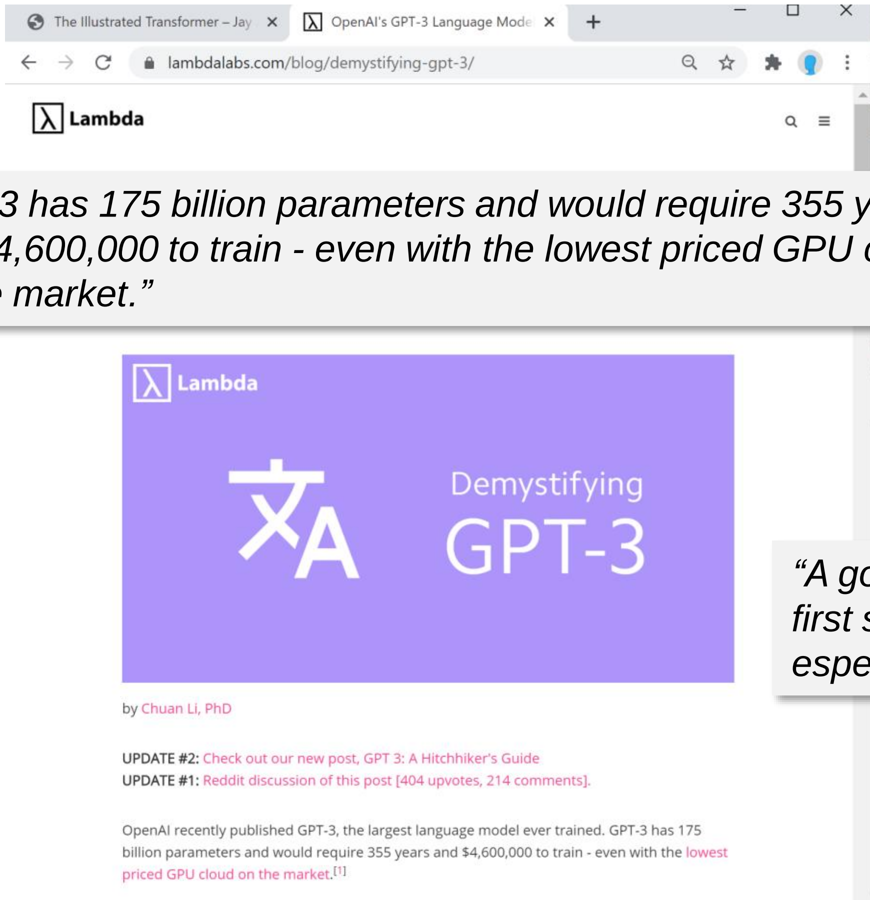

"GPT-3 has 175 billion parameters and would require 355 years and $4,600,000 to train - even with the lowest priced GPU cloud on the market."

vs.

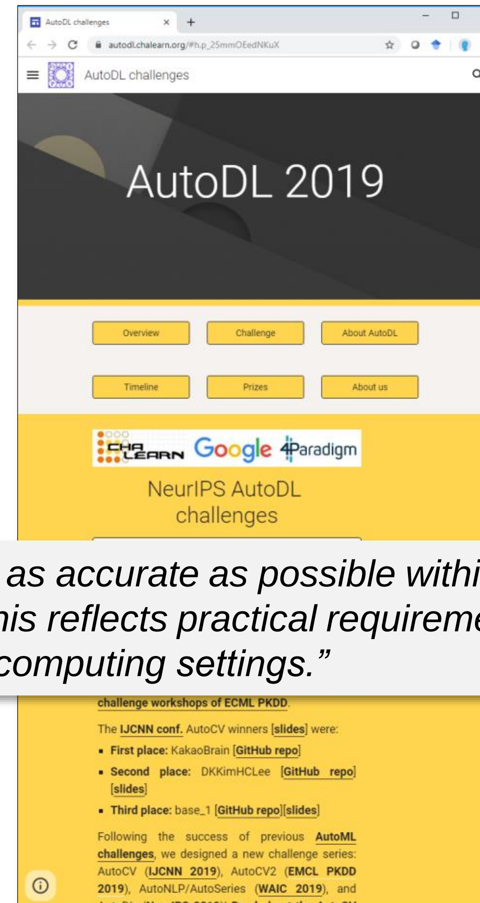

"A good system thus has to be as accurate as possible within the first seconds of training […]. This reflects practical requirements especially in mobile and edge computing settings."

Sources: https://lambdalabs.com/blog/demystifying-gpt-3/, Tuggener et al. (2929), *«Design Patterns for Resource Constrained Automated Deep Learning Methods»*, submitted to MDPI AI

The problem of deployment
Introducing MLOps

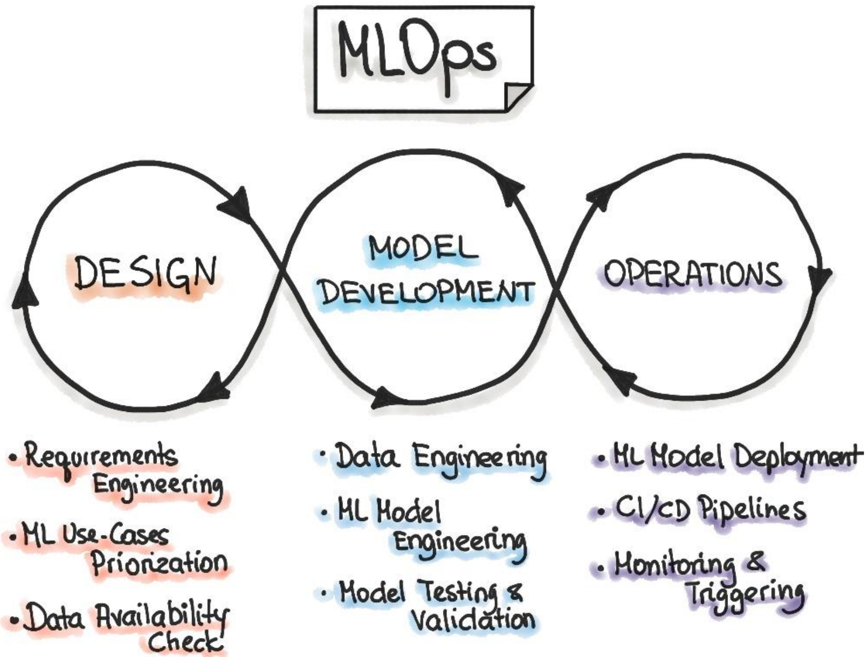

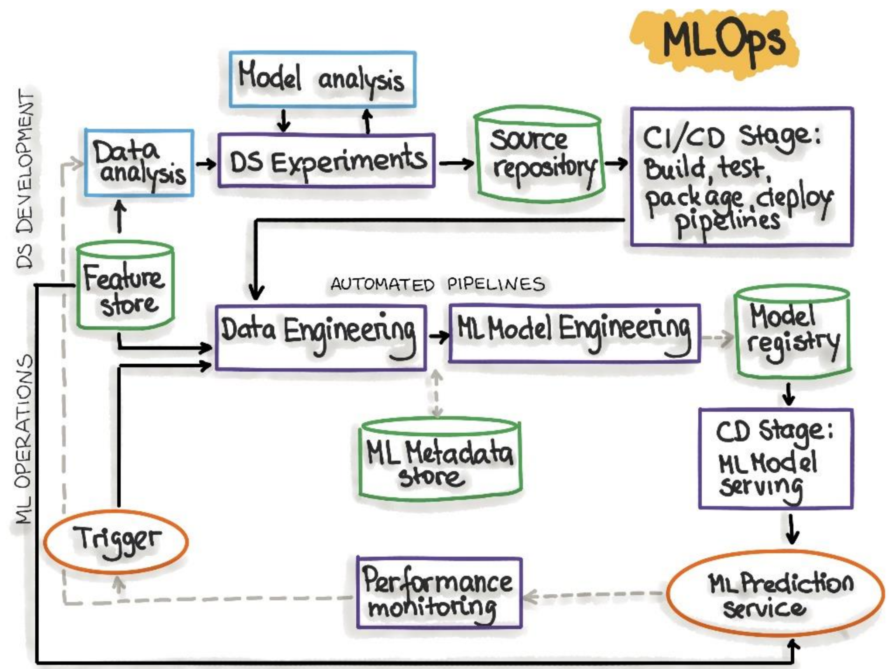

Source: INNOQ / https://ml-ops.org/content/mlops-principles

Zürcher Hochschule
für Angewandte Wissenschaften

Conclusions

- Deep learning **is applied** and deployed in «normal» businesses (non-AI, SME)
- It does not need big-, but some **data** (**effort usually underestimated**)
- IT needs to consider: specific **risks in procurement** / customization, **computational resources**, continued development after deployment (**MLOps**),

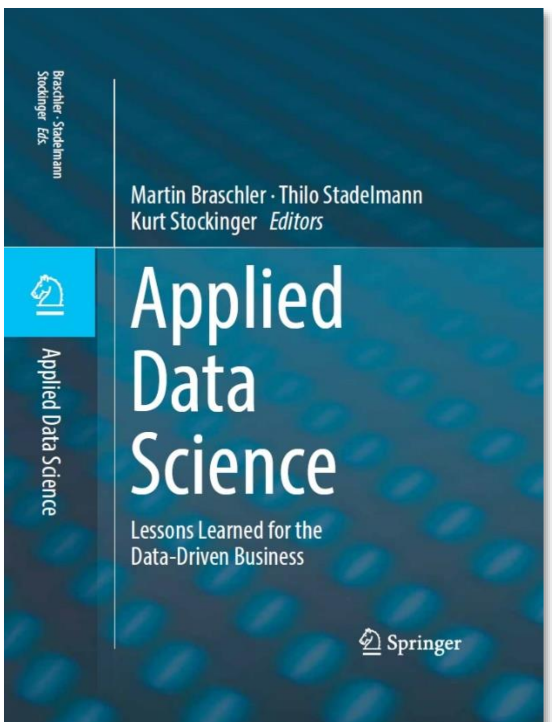

About me:

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

Further contacts:

- datalab@zhaw.ch, info.office@data-service-alliance.ch, office-switzerland@claire-ai.org

➔ Happy to answer questions & requests

datalab
www.zhaw.ch/datalab

data service
Swiss Alliance for Data-Intensive Services

CLAIRE

APPENDIX

Zürcher Hochschule
für Angewandte Wissenschaften

zh
aw

About us & our work

ZHAW School of Engineering, Winterthur, Switzerland

Machine learning-based
Pattern Recognition

- Robust applications
- Biometrics
- Document Analysis
- Learning to act

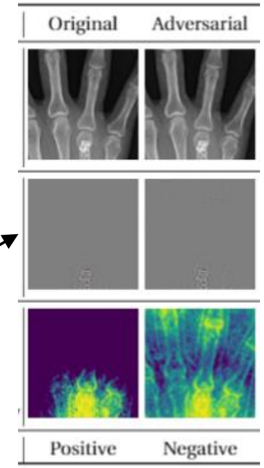

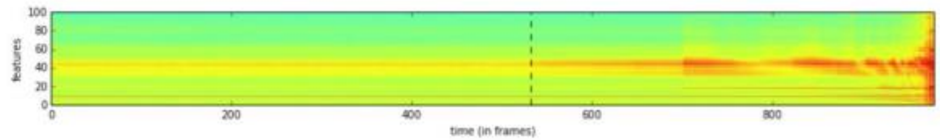

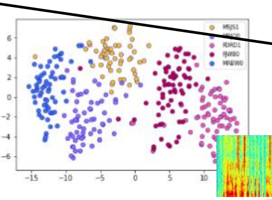

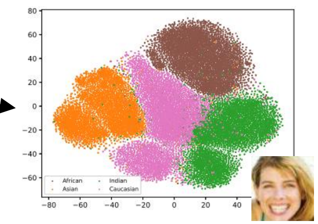

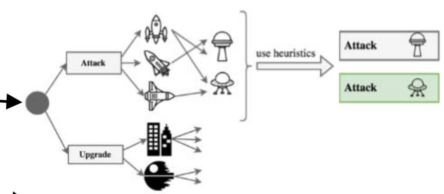

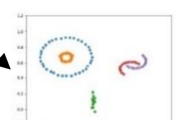

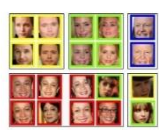

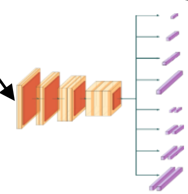

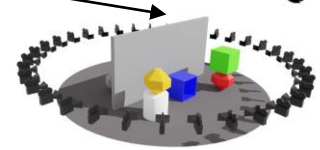

Outlook: recent work in progress

- Learning to reduce motion artifacts in 3D CT scans

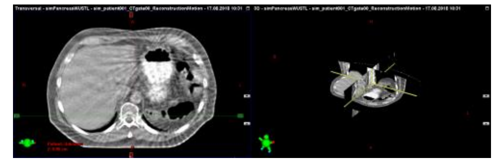

- Learning an artificial communication language for multi-agent reinforcement learning in logistics
  (notable rank in Flatland 2019 competition)

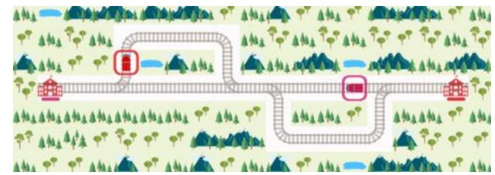

- Automated deep learning
  (top rank in AutoDL 2020 challenge)

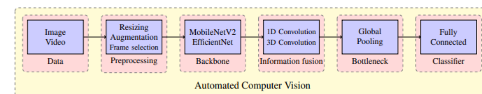

- Learning to segment and classify food waste in professional kitchens under adversarial conditions

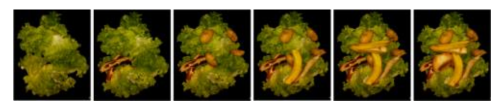

- Improving robotic vision through active vision and combined supervised and reinforcement learning
  (Dr. Waldemar Jucker Award 2020)

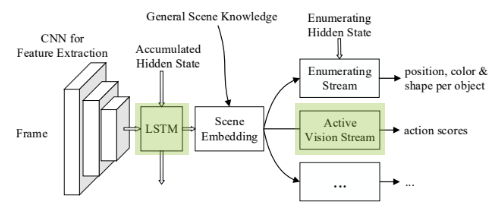

Roost, Meier, Huschauer, Nygren, Egli, Weiler & Stadelmann (2020). *«Improving Sample Efficiency and Multi-Agent Communication in RL-based Train Rescheduling»*. SDS'2020.
Tuggener, Amirian, Benites, von Däniken, Gupta, Schilling & Stadelmann (2020). *«Design Patterns for Resource Constrained Automated Deep Learning Methods».* Submitted to MDPI AI.
Roost, Meier, Toffetti Carughi & Stadelmann (2020). *«Combining Reinforcement Learning with Supervised Deep Learning for Neural Active Scene Understanding».* AVHRC 2020

Foundation
Inductive supervised learning

Assumption
- A model fitted to a *sufficiently large* sample of data…
- …will **generalize** to unseen data

Method
- **Searching for optimal parameters of a function**…
- …such that all sample inputs (images) are mapped to the correct outputs (e.g., «car»)

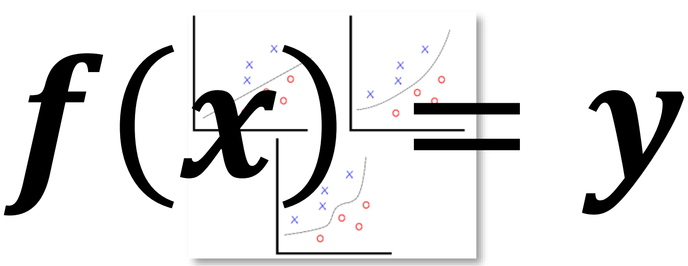

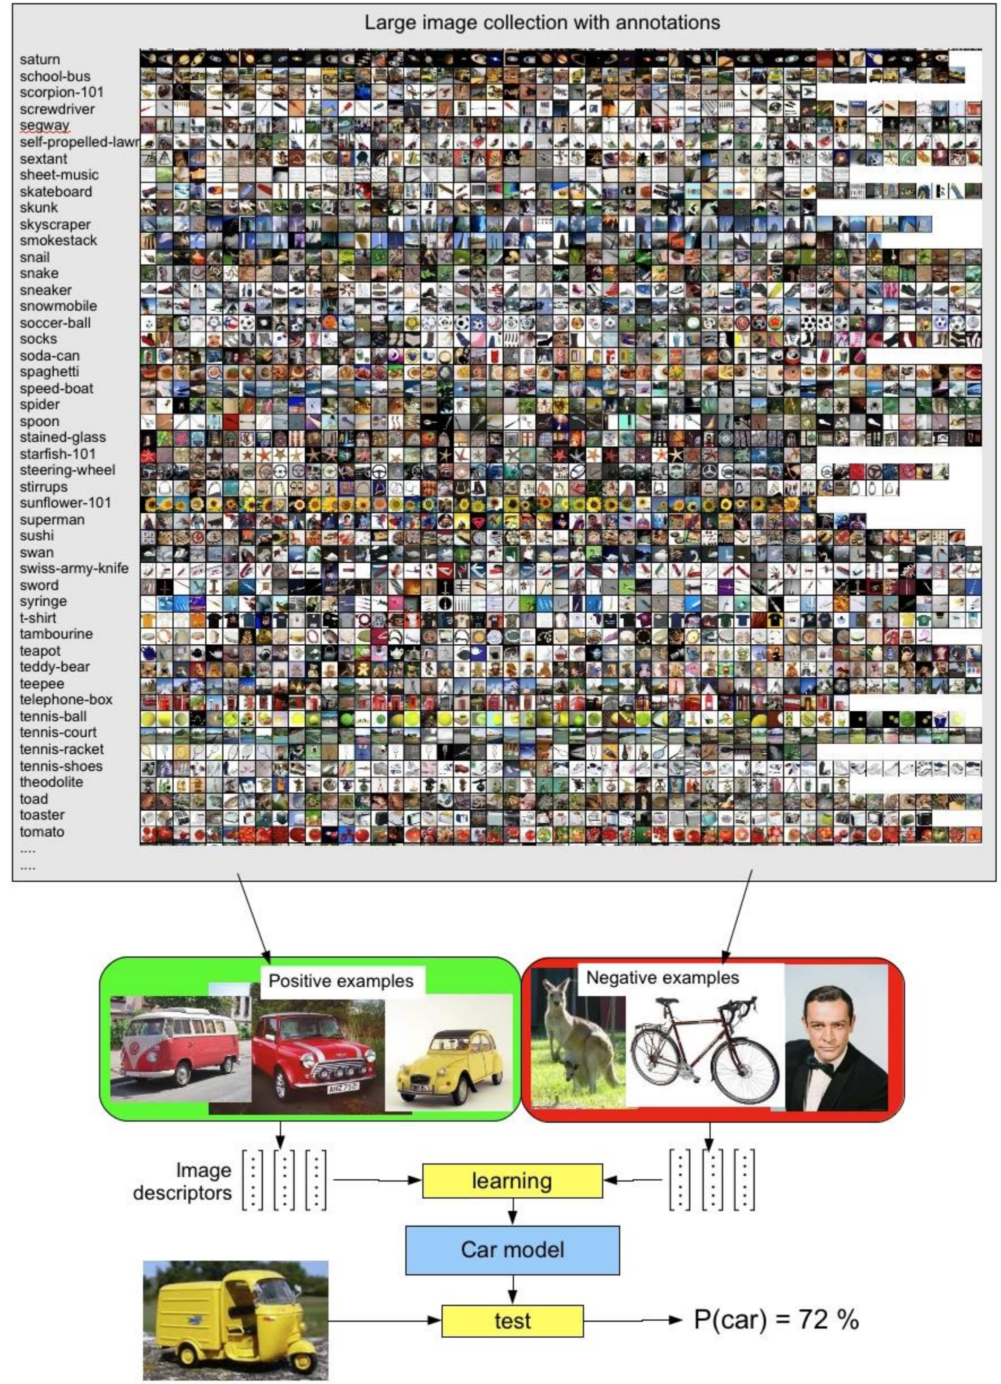

Quelle: http://lear.inrialpes.fr/job/postdoc-large-scale-classif-11-img/attribs_patchwork.jpg

Search for optimal parameters *of a function*?

neuron

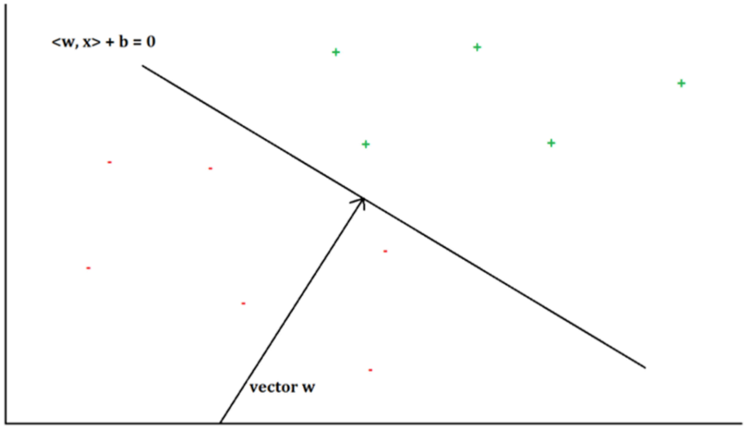

neural net

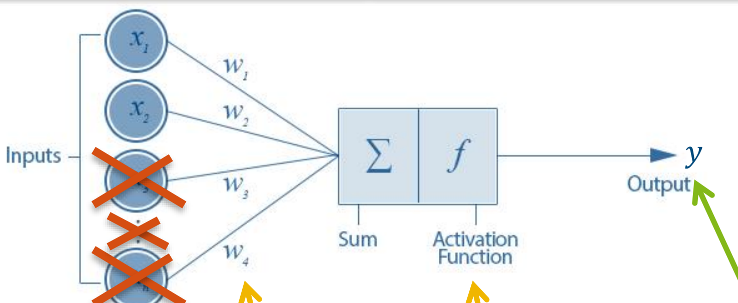

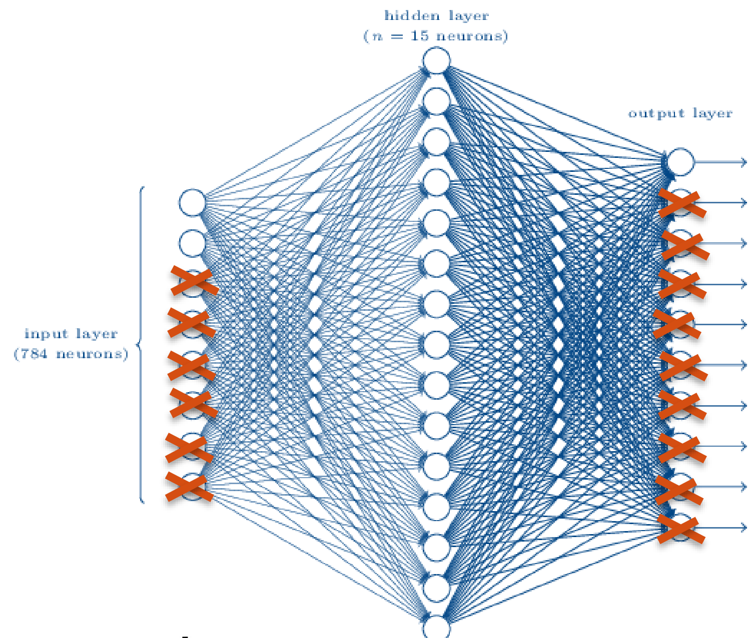

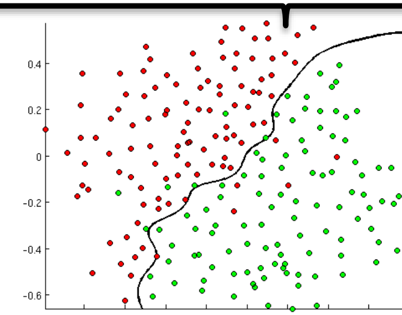

Threshold / decision

features (e.g., pixel)

Result (interpreted as e.g. «car»)

Tunable parameters

Search for optimal parameters of a function?

Probability [%] for specific event

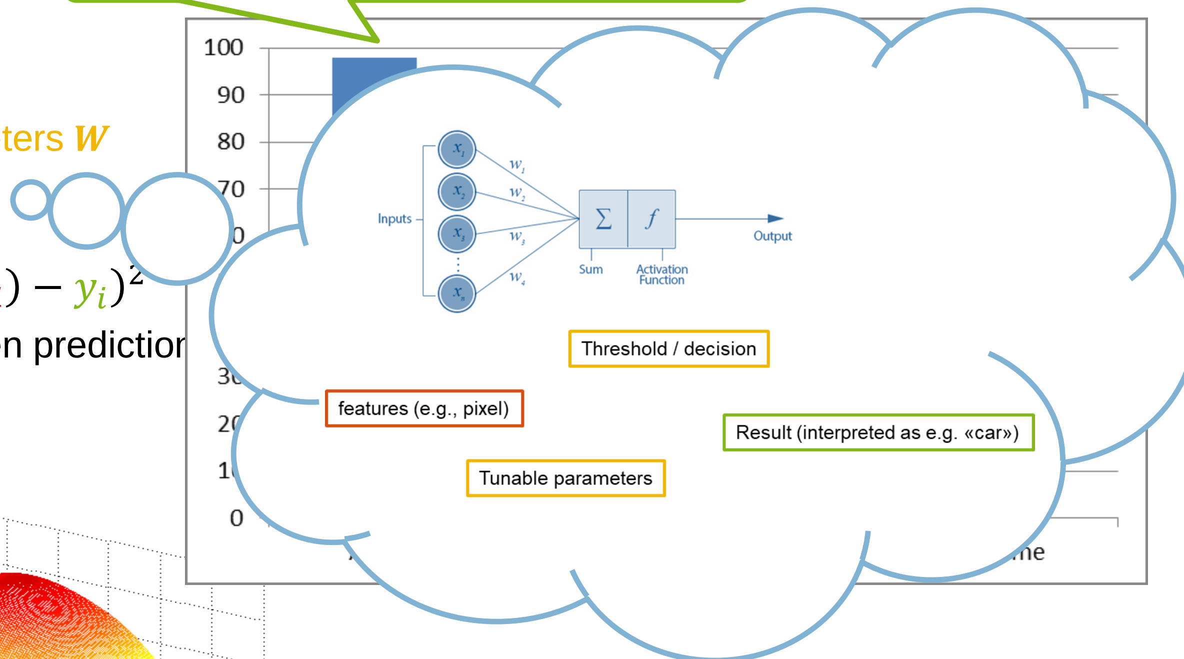

- Our artificial neural net: $f_W(x) = y$
  with image x, ground truth y and parameters $\boldsymbol{W}$
  ($\boldsymbol{W} = \{w_1, w_2\}$ initialized at random)

- Error measure: $l(W) = \frac{1}{N}\sum_{i=1}^{N}(f_W(x_i) - y_i)^2$
  Average of (quadratic) difference between prediction and ground truth («loss»)

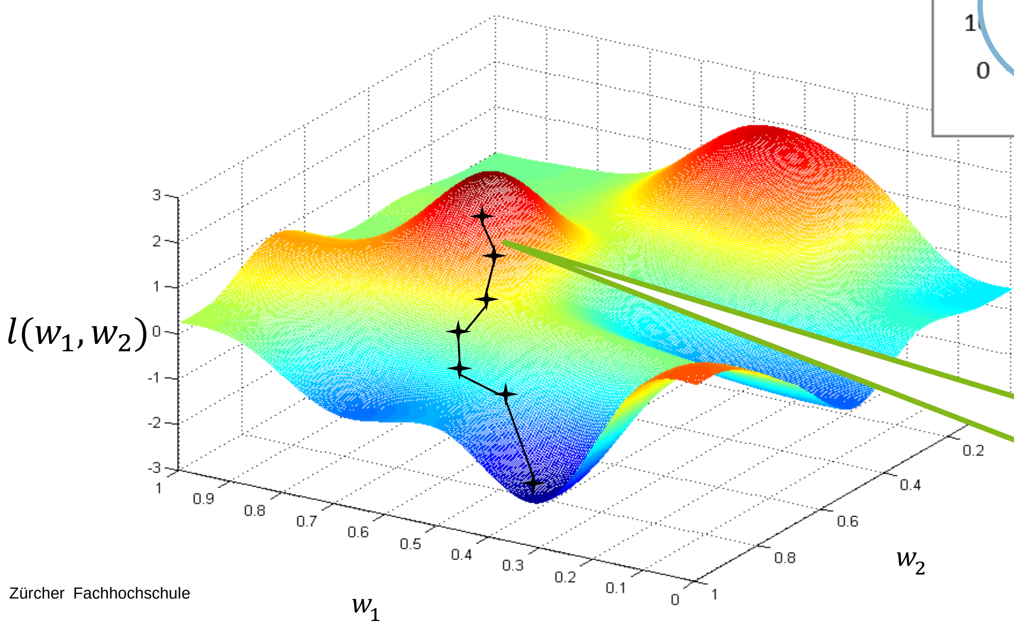

← error landscape

Method: iterative change of parameters of f in the direction of the steepest descent of J

What does the neural network «see»?
Hierarchy of more complex features

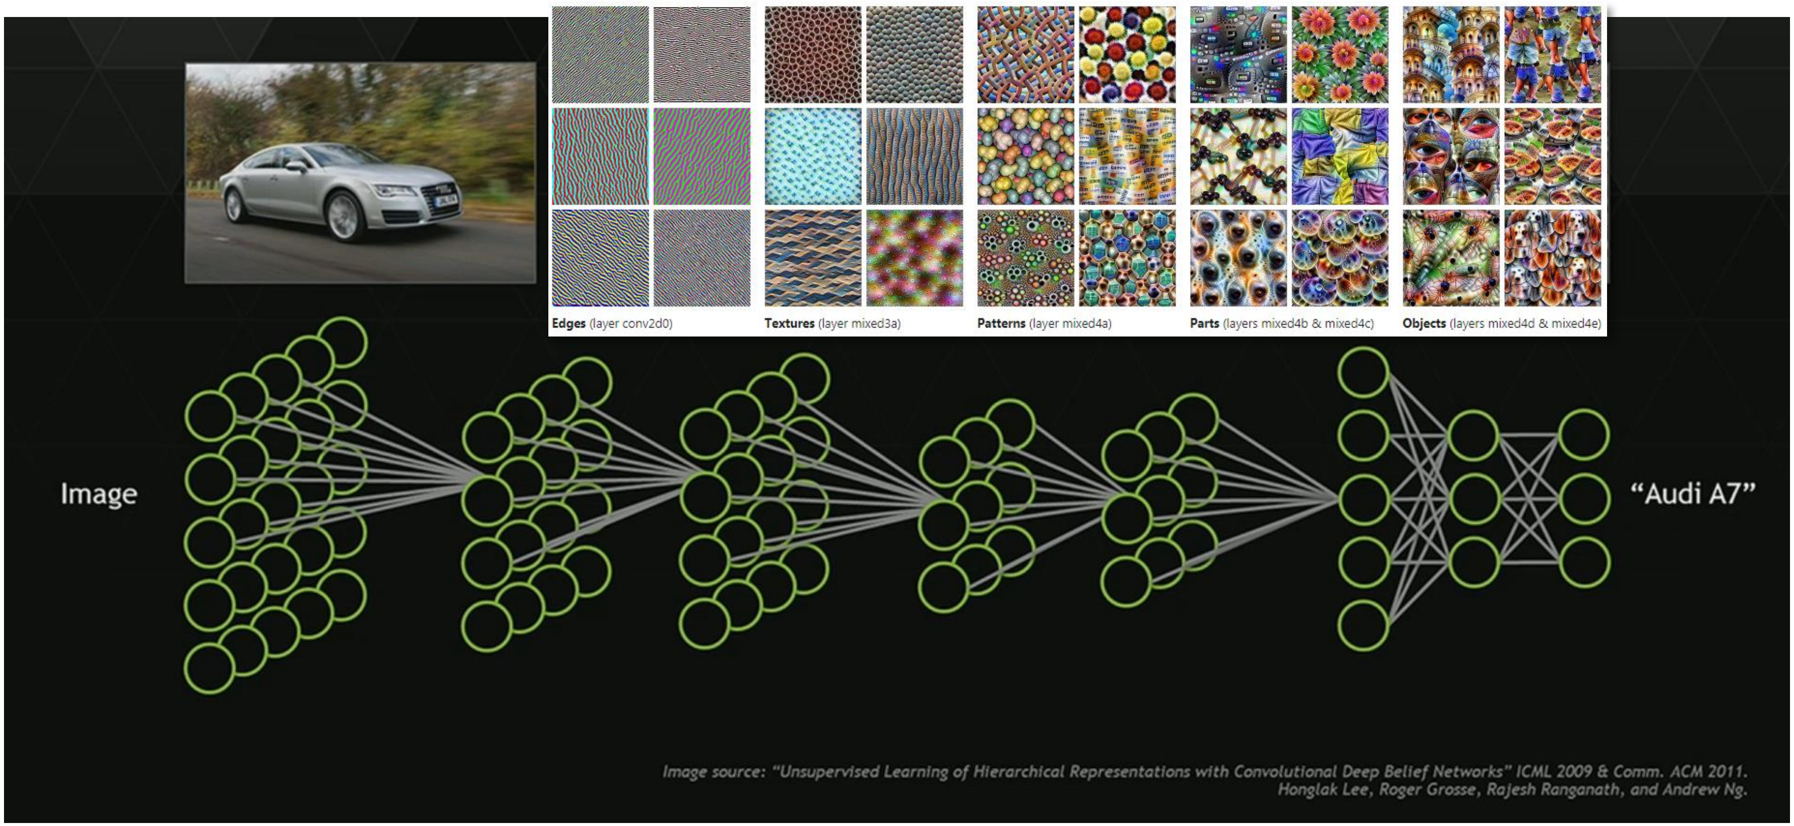

Source: https://www.pinterest.com/explore/artificial-neural-network/
Olah, et al., "Feature Visualization", Distill, 2017, https://distill.pub/2017/feature-visualization/.