From Loss to Love:
Lessons Learned in Deep Learning Research

Hexagon Technology Forum

Prof. Dr. Thilo Stadelmann

ZHAW School of Engineering

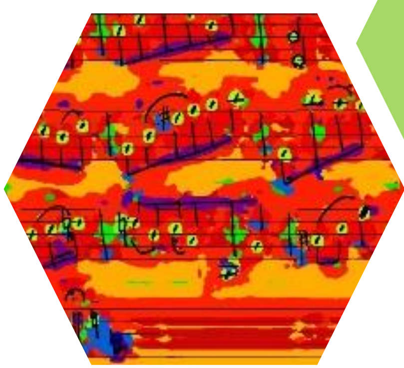

Deep learning has matured considerably in recent years. This talk explores current results from applied deep learning research. Shown examples include object detection, industrial quality control, predictive maintenance, document analysis and speech processing. We report on their impact on daily work by formulating specific lessons learned and thereby connect the dots from hands-on development work (loss) to societal benefit (love).

Heerbrugg

Multimediaraum
09:00am – 11:00am

May 23, 2019

HEXAGON

From Loss to Love: Lessons Learned in Deep Learning Research

Hexagon Technology Forum

Agenda

Welcome	Pascal Jordil	*5 min*
A brief Intro to AI, ML & DL	Bernd Reimann	*10 min*
From Loss to Love: Lessons Learned in Deep Learning Research	Thilo Stadelmann	*75 min*
AI R&D within Hexagon	Bernd Reimann	*10 min*
Open Discussion	all	*20 min*

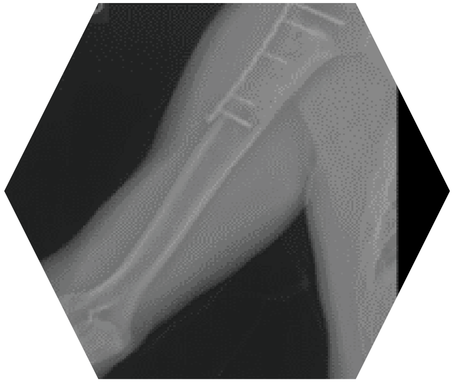

HEXAGON

Zürcher Hochschule
für Angewandte Wissenschaften

From Loss to Love:
Lessons Learned in Deep Learning Research

Hexagon Technology Forum, CH-Heerbrugg, May 23, 2019

Thilo Stadelmann

Zürcher Hochschule
für Angewandte Wissenschaften

About us

ML @ Information Engineering Group, ZHAW InIT






Machine learning-based Pattern Recognition

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


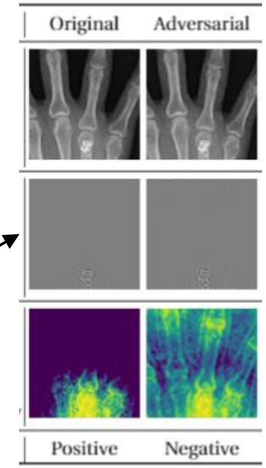

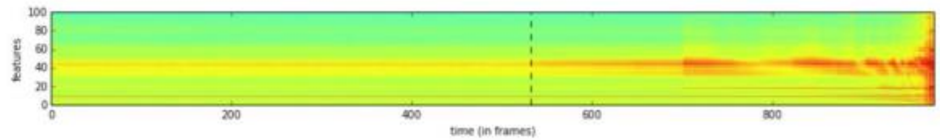

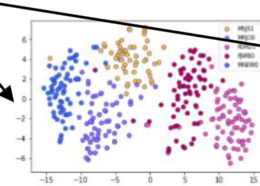
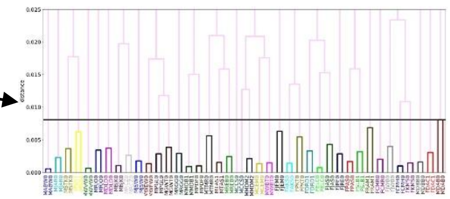



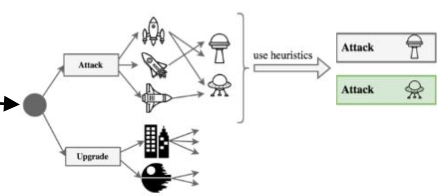

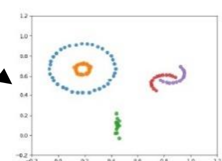
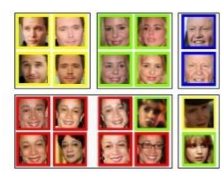

Why deep learning?

arXiv monthly submission rates

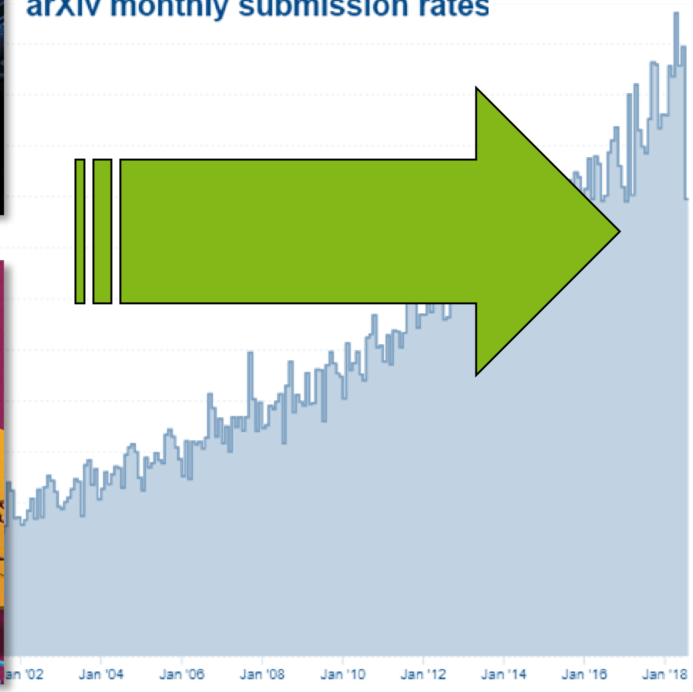

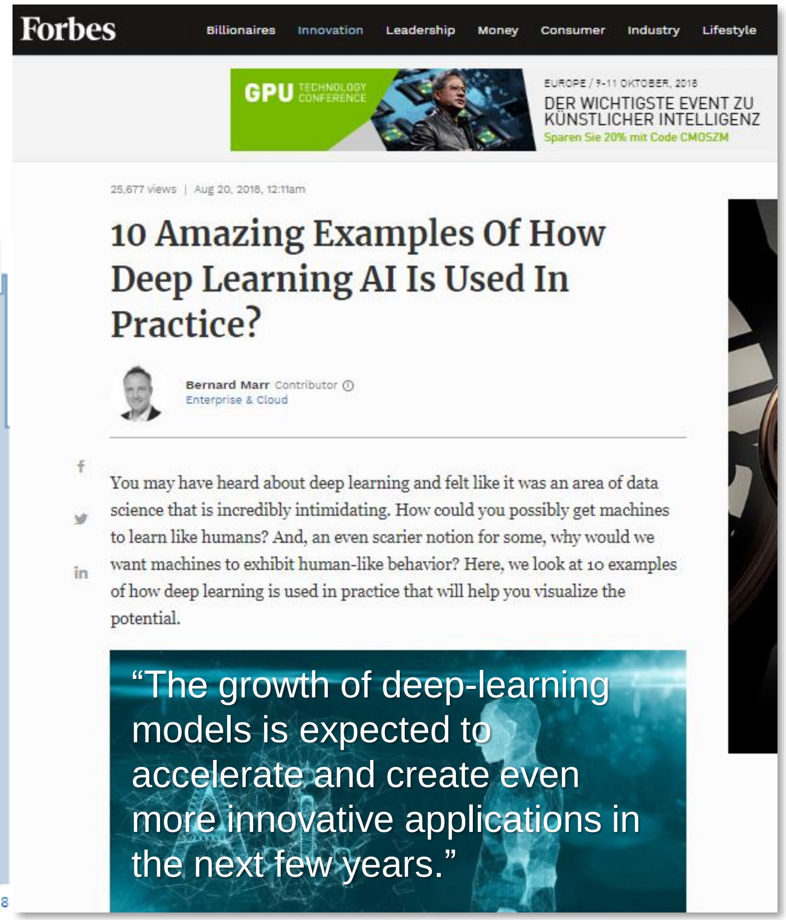

Forbes Billionaires Innovation Leadership Money Consumer Industry Lifestyle

GPU TECHNOLOGY CONFERENCE EUROPE / 9-11 OKTOBER, 2018 DER WICHTIGSTE EVENT ZU KÜNSTLICHER INTELLIGENZ Sparen Sie 20% mit Code CMOSZM

25,677 views | Aug 20, 2018, 12:11am

10 Amazing Examples Of How Deep Learning AI Is Used In Practice?

Bernard Marr Contributor
Enterprise & Cloud

You may have heard about deep learning and felt like it was an area of data science that is incredibly intimidating. How could you possibly get machines to learn like humans? And, an even scarier notion for some, why would we want machines to exhibit human-like behavior? Here, we look at 10 examples of how deep learning is used in practice that will help you visualize the potential.

"The growth of deep-learning models is expected to accelerate and create even more innovative applications in the next few years."

Zürcher Hochschule
für Angewandte Wissenschaften

Why deep learning?

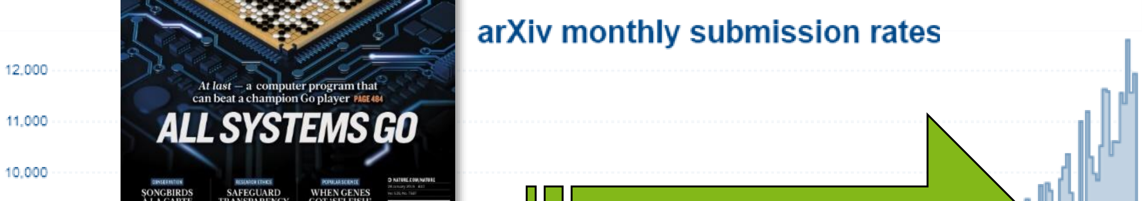

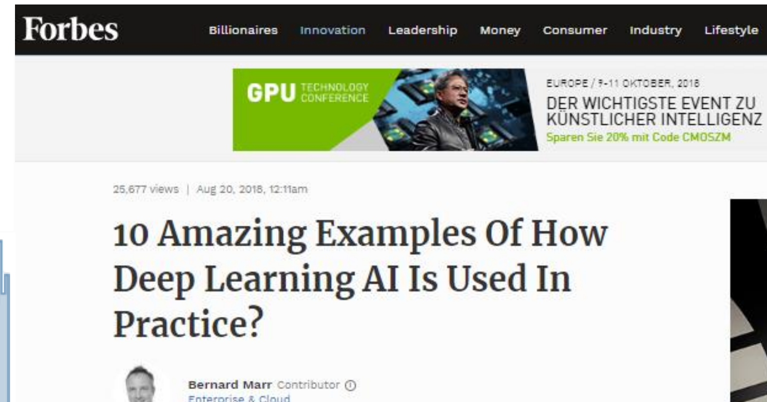

25,677 views | Aug 20, 2018, 12:11am

10 Amazing Examples Of How Deep Learning AI Is Used In Practice?

Bernard Marr Contributor
Enterprise & Cloud

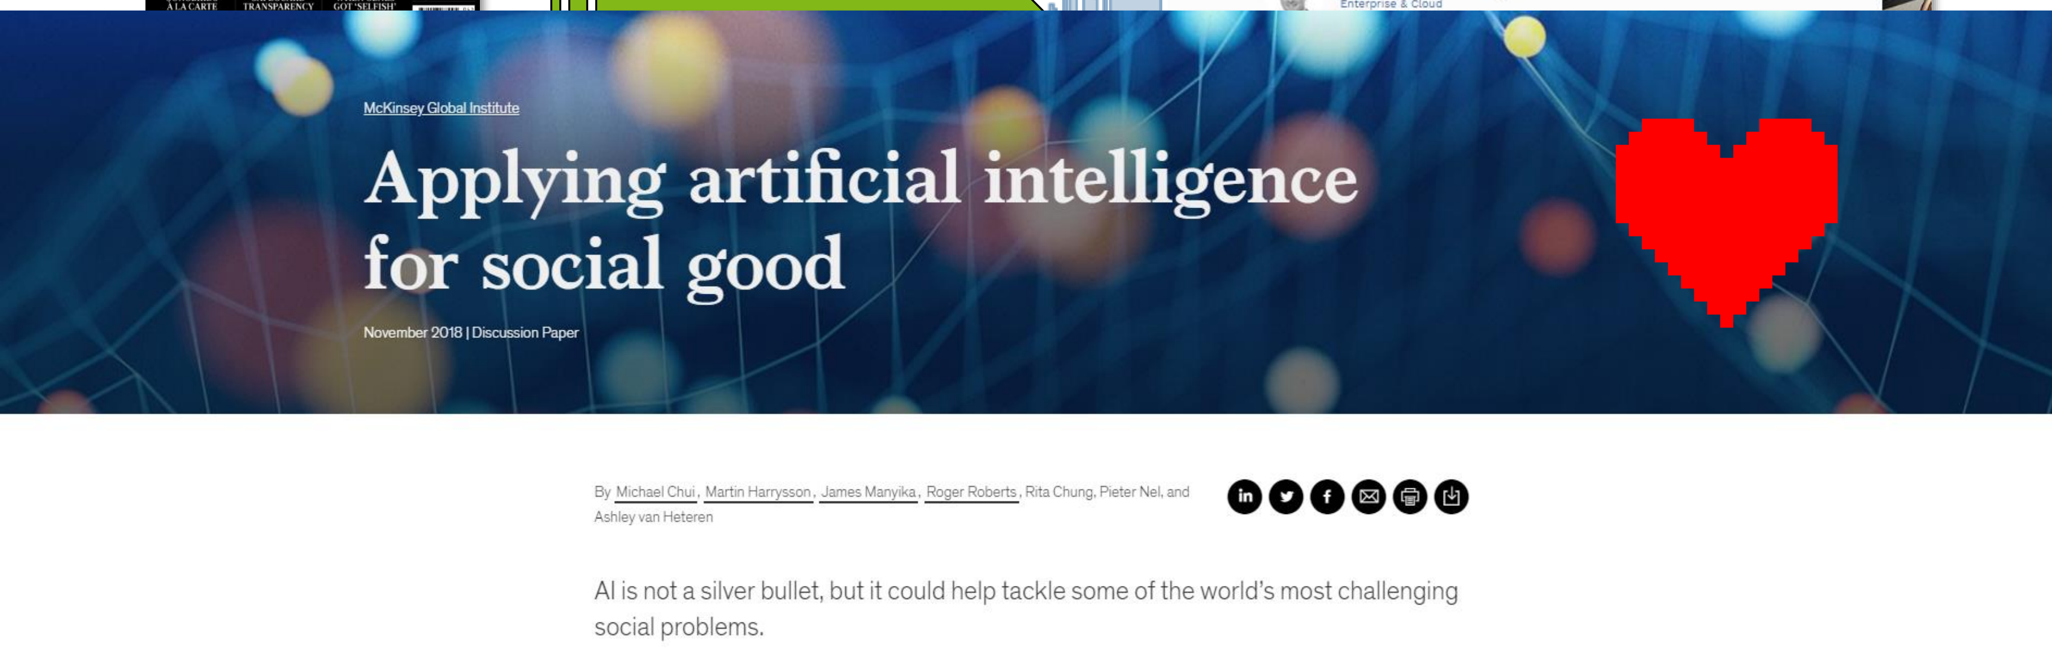

McKinsey Global Institute

Applying artificial intelligence for social good

November 2018 | Discussion Paper

By Michael Chui, Martin Harrysson, James Manyika, Roger Roberts, Rita Chung, Pieter Nel, and Ashley van Heteren

AI is not a silver bullet, but it could help tackle some of the world's most challenging social problems.

Idea: Add depth to learn features automatically

Classical image processing

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

(0.2, 0.4, …)

(0.4, 0.3, …)

Classification
(SVM, neural network, etc.)

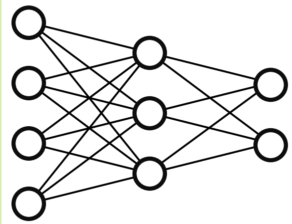

Container ship

Tiger

…

Idea: Add depth to learn features automatically

Classical image processing

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

(0.2, 0.4, …)

(0.4, 0.3, …)

Classification (SVM, neural network, etc.)

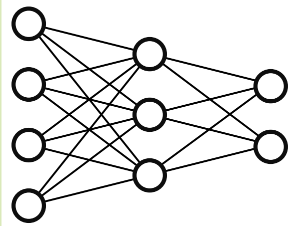

Container ship

Tiger

…

Using Convolutional Neual Networks (CNNs)

Takes raw pixels in, learns features automatically!

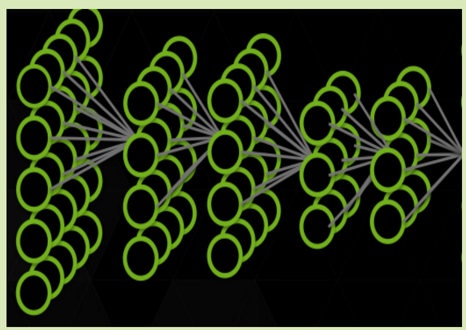

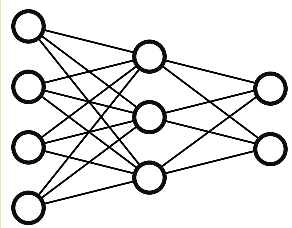

Container ship

Tiger

…

Idea: Add depth to learn features automatically

Classical image processing

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

(0.2, 0.4, …)

Classification
(SVM, neural network, etc.)

Container ship

Tiger

…

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

Using Convolutional Neual Networks (CNNs)

Takes raw pixels in, learns features automatically!

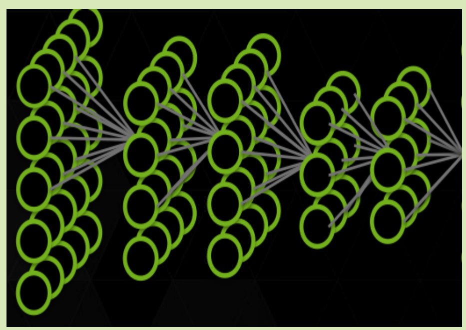

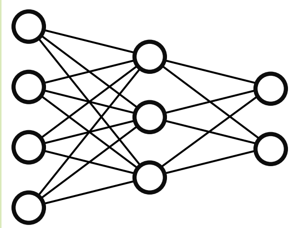

Container ship

Tiger

…

Agenda

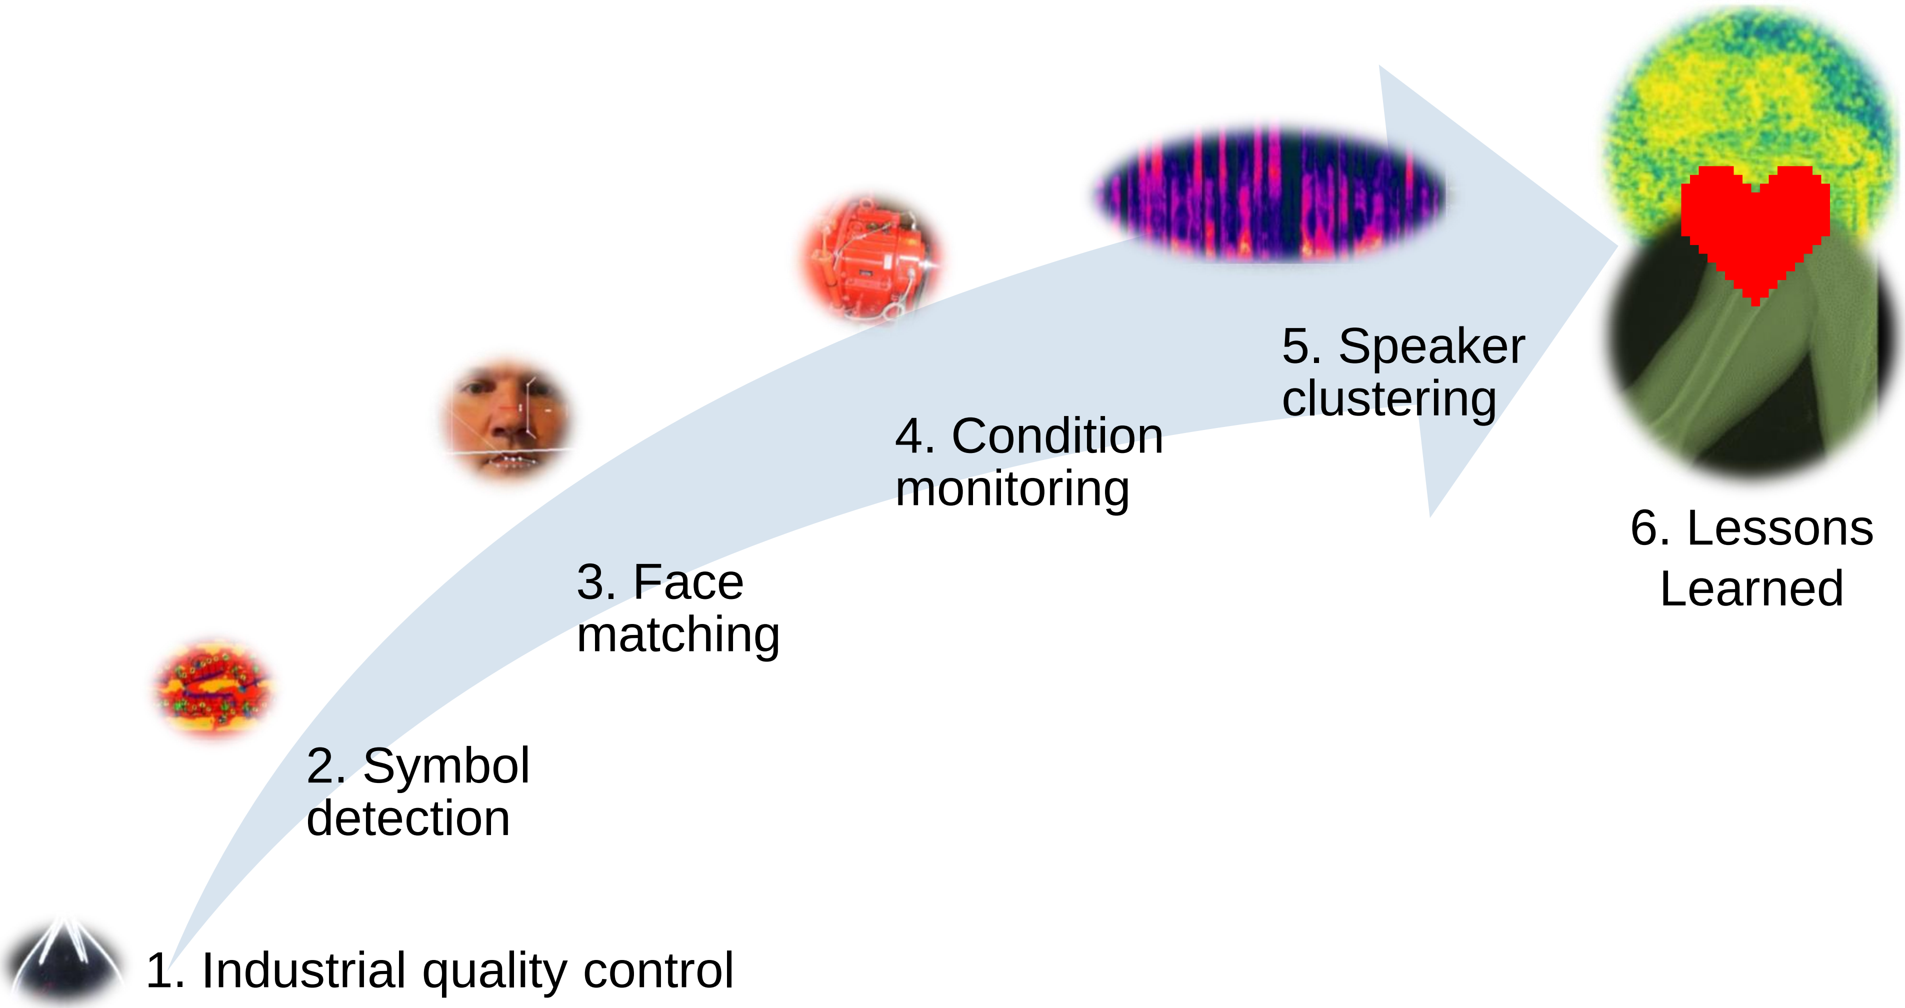
1. Industrial quality control
2. Symbol detection
3. Face matching
4. Condition monitoring
5. Speaker clustering
6. Lessons Learned

1. Industrial quality control

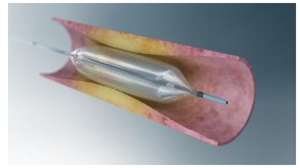

Zürcher Hochschule
für Angewandte Wissenschaften

zh
aw

BW-TEC
ENGINEERING AUTOMATION CONSULTING

Schweizerische Eidgenossenschaft
Confédération suisse
Confederazione Svizzera
Confederaziun svizra

Swiss Confederation

Innosuisse – Swiss Innovation Agency

Task

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

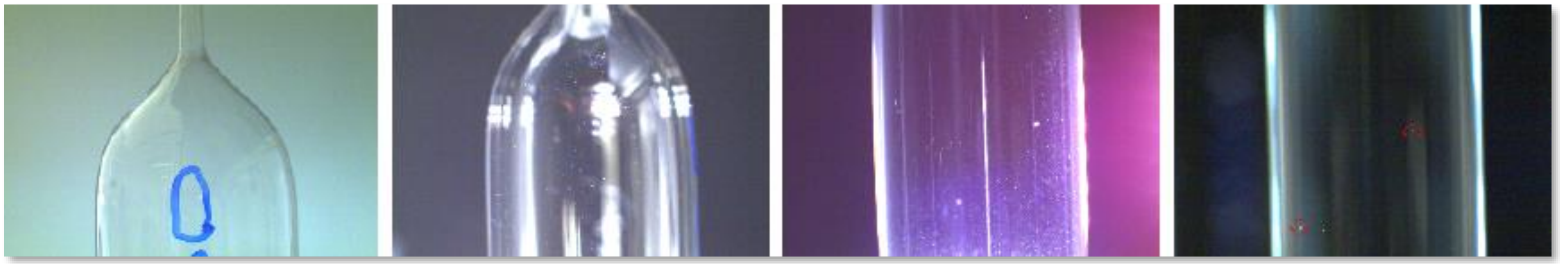

Challenges

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

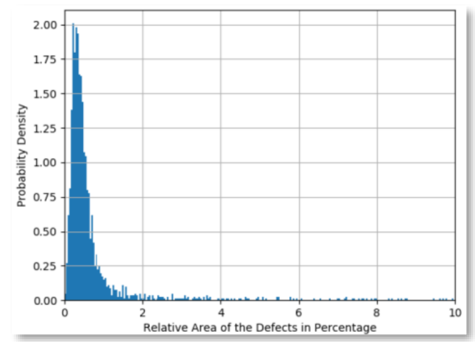

Stadelmann, Amirian, Arabaci, Arnold, Duivesteijn, Elezi, Geiger, Lörwald, Meier, Rombach & Tuggener (2018). *«Deep Learning in the Wild»*. ANNPR’2018.

1. Industrial quality control – baseline results

Zürcher Hochschule
für Angewandte Wissenschaften

Ingredients
- Weighted loss
- Defect cropping
- Careful customization

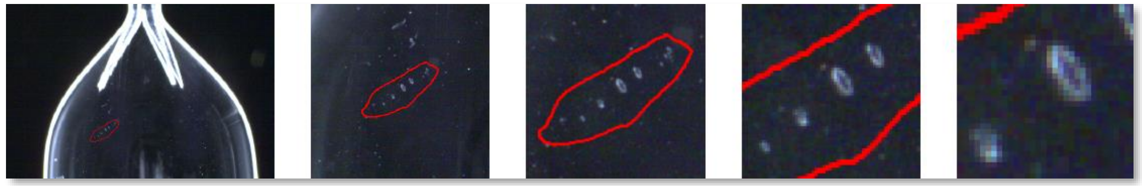

Schweizerische Eidgenossenschaft
Confédération suisse
Confederazione Svizzera
Confederaziun svizra

Swiss Confederation

Innosuisse – Swiss Innovation Agency

Interim results

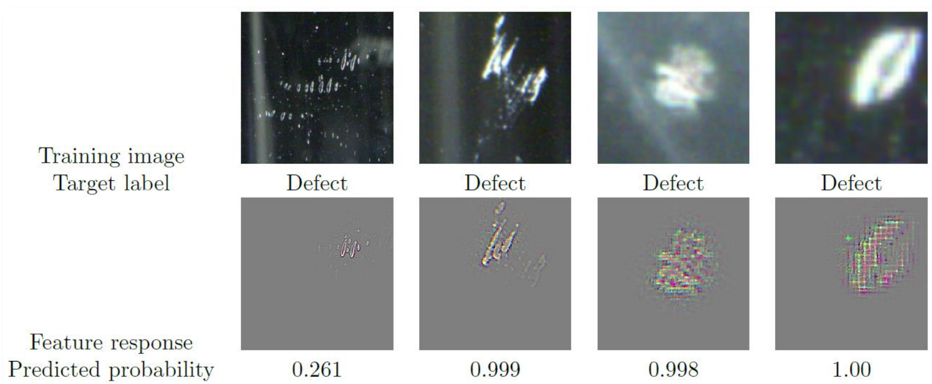

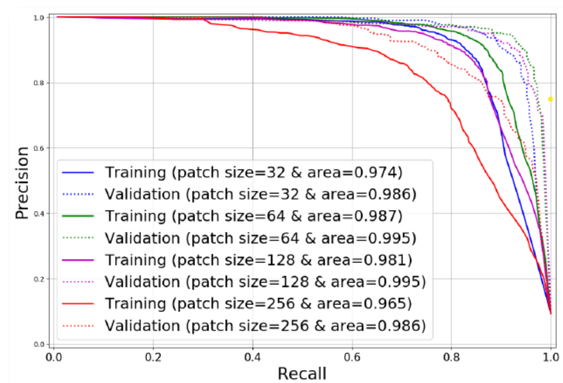

Zürcher Hochschule
für Angewandte Wissenschaften

1. Industrial quality control – recent results

- Human performance isn't flawless
- Tailoring pays off

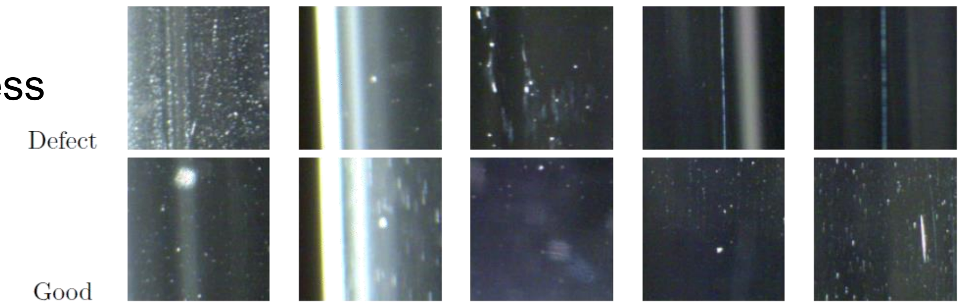

Figure 2: Samples of failure cases in classification. The shown *defect* samples in the table are not recognized as a defects, and the *good* images are misclassified as defects.

Schweizerische Eidgenossenschaft
Confédération suisse
Confederazione Svizzera
Confederaziun svizra

Swiss Confederation

Innosuisse – Swiss Innovation Agency

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

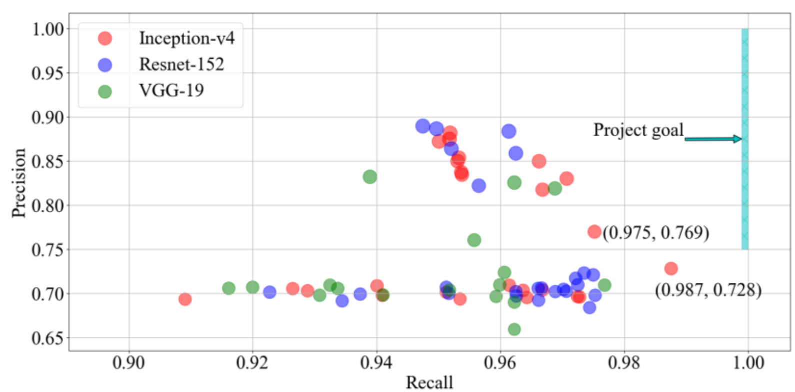

1. Industrial quality control – future work?

Schweizerische Eidgenossenschaft
Confédération suisse
Confederazione Svizzera
Confederaziun svizra

Swiss Confederation

Innosuisse – Swiss Innovation Agency

Approaches to overcome class imbalance and small training set sizes?

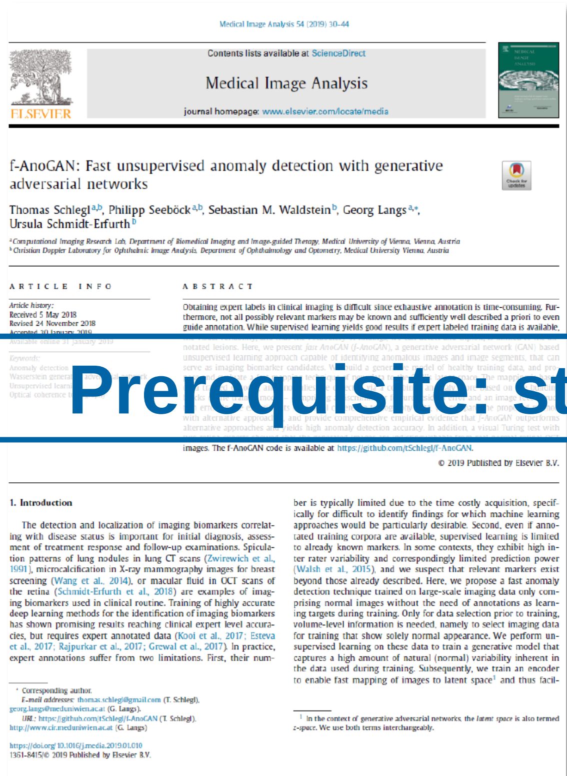

Medical Image Analysis 54 (2019) 30–44

Contents lists available at ScienceDirect

Medical Image Analysis

journal homepage: www.elsevier.com/locate/media

f-AnoGAN: Fast unsupervised anomaly detection with generative adversarial networks

Thomas Schlegl[a,b], Philipp Seeböck[a,b], Sebastian M. Waldstein[b], Georg Langs[a,*], Ursula Schmidt-Erfurth[b]

[a] Computational Imaging Research Lab, Department of Biomedical Imaging and Image-guided Therapy, Medical University of Vienna, Vienna, Austria
[b] Christian Doppler Laboratory for Ophthalmic Image Analysis, Department of Ophthalmology and Optometry, Medical University Vienna, Austria

ARTICLE INFO

Article history:
Received 5 May 2018
Revised 24 November 2018
Accepted 30 January 2019
Available online 31 January 2019

Keywords:
Anomaly detection

ABSTRACT

Obtaining expert labels in clinical imaging is difficult since exhaustive annotation is time-consuming. Furthermore, not all possibly relevant markers may be known and sufficiently well described a priori to even guide annotation. While supervised learning yields good results if expert labeled training data is available,

notated lesions. Here, we present *fast AnoGAN* (*f-AnoGAN*), a generative adversarial network (GAN) based unsupervised learning approach capable of identifying anomalous images and image segments, that can serve as imaging biomarker candidates.

with alternative approaches, and provide comprehensive empirical evidence that *f-AnoGAN* outperforms alternative approaches and yields high anomaly detection accuracy. In addition, a visual Turing test with

images. The f-AnoGAN code is available at https://github.com/tSchlegl/f-AnoGAN.

1. Introduction

The detection and localization of imaging biomarkers correlating with disease status is important for initial diagnosis, assessment of treatment response and follow-up examinations. Spiculation patterns of lung nodules in lung CT scans (Zwirewich et al., 1991), microcalcification in X-ray mammography images for breast screening (Wang et al., 2014), or macular fluid in OCT scans of the retina (Schmidt-Erfurth et al., 2018) are examples of imaging biomarkers used in clinical routine. Training of highly accurate deep learning methods for the identification of imaging biomarkers has shown promising results reaching clinical expert level accuracies, but requires expert annotated data (Kooi et al., 2017; Esteva et al., 2017; Rajpurkar et al., 2017; Grewal et al., 2017). In practice, expert annotations suffer from two limitations. First, their number is typically limited due to the time costly acquisition, specifically for difficult to identify findings for which machine learning approaches would be particularly desirable. Second, even if annotated training corpora are available, supervised learning is limited to already known markers. In some contexts, they exhibit high inter rater variability and correspondingly limited prediction power (Walsh et al., 2015), and we suspect that relevant markers exist beyond those already described. Here, we propose a fast anomaly detection technique trained on large-scale imaging data only comprising normal images without the need of annotations as learning targets during training. Only for data selection prior to training, volume-level information is needed, namely to select imaging data for training that show solely normal appearance. We perform unsupervised learning on these data to train a generative model that captures a high amount of natural (normal) variability inherent in the data used during training. Subsequently, we train an encoder to enable fast mapping of images to latent space[1] and thus facil-

* Corresponding author.
E-mail addresses: thomas.schlegl@gmail.com (T. Schlegl), georg.langs@meduniwien.ac.at (G. Langs).
URL: https://github.com/tSchlegl/f-AnoGAN (T. Schlegl), http://www.cir.meduniwien.ac.at (G. Langs)

[1] In the context of generative adversarial networks, the *latent space* is also termed *z-space*. We use both terms interchangeably.

https://doi.org/10.1016/j.media.2019.01.010
1361-8415/© 2019 Published by Elsevier B.V.

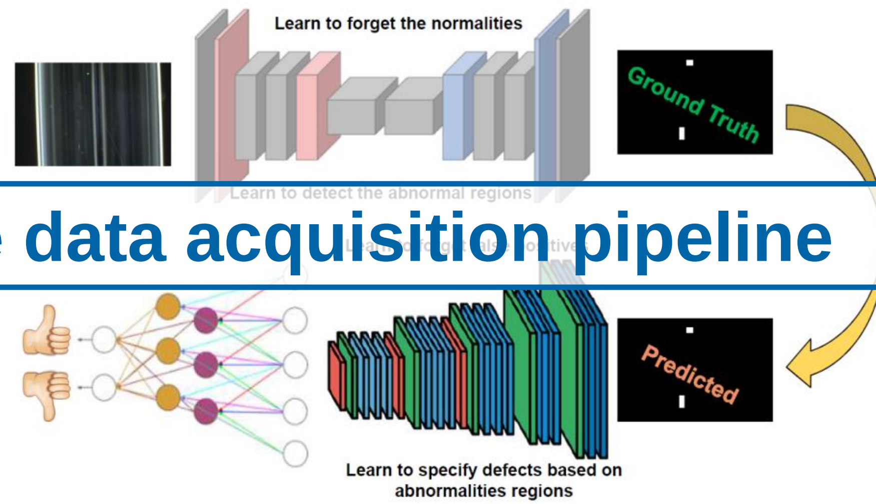

Prerequisite: stable data acquisition pipeline

2. Symbol detection

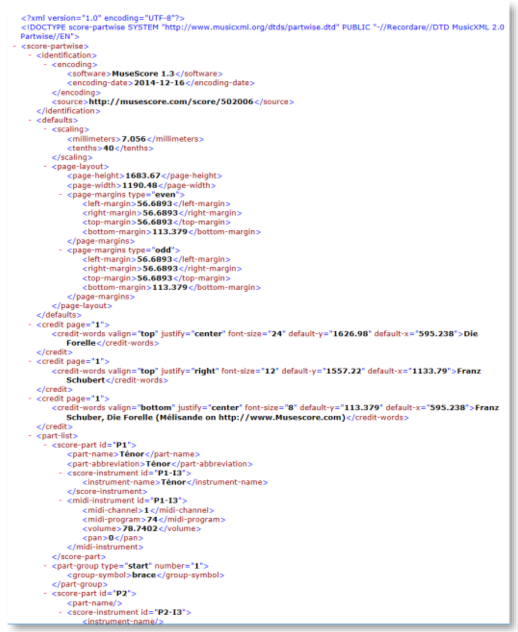

```
<?xml version="1.0" encoding="UTF-8"?>
<!DOCTYPE score-partwise SYSTEM "http://www.musicxml.org/dtds/partwise.dtd" PUBLIC "-//Recordare//DTD MusicXML 2.0 Partwise//EN">
- <score-partwise>
  - <identification>
    - <encoding>
        <software>MuseScore 1.3</software>
        <encoding-date>2014-12-16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


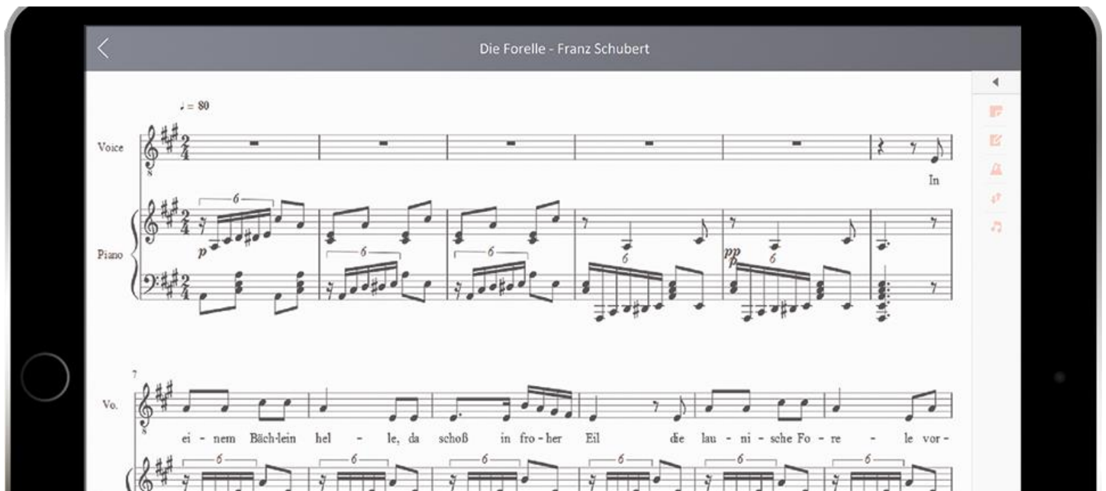

Zürcher Hochschule
für Angewandte Wissenschaften

SCOREPAD

Schweizerische Eidgenossenschaft
Confédération suisse
Confederazione Svizzera
Confederaziun svizra

Swiss Confederation

Innosuisse – Swiss Innovation Agency

2. Symbol detection – challenges & solutions

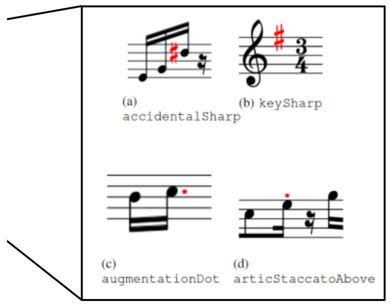

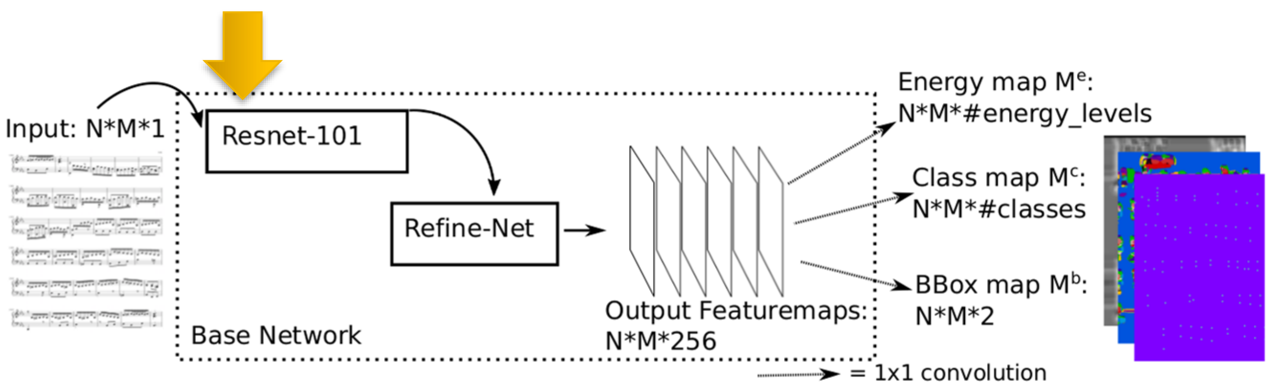

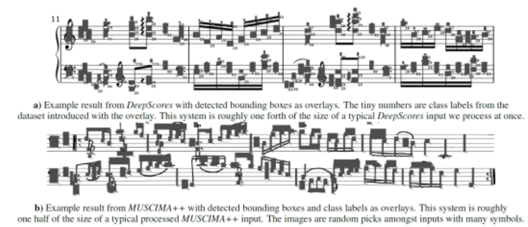

Tuggener, Elezi, Schmidhuber, Pelillo & Stadelmann (2018). *«DeepScores – A Dataset for Segmentation, Detection and Classification of Tiny Objects»*. ICPR'2018.
Tuggener, Elezi, Schmidhuber & Stadelmann (2018). *«Deep Watershed Detector for Music Object Recognition»*. ISMIR'2018.

Zürcher Hochschule
für Angewandte Wissenschaften

2. Symbol detection – methodology
OMR vs state of the art object detectors

SCOREPAD

Schweizerische Eidgenossenschaft
Confédération suisse
Confederazione Svizzera
Confederaziun svizra

Swiss Confederation

Innosuisse – Swiss Innovation Agency

YOLO/SSD-type detectors

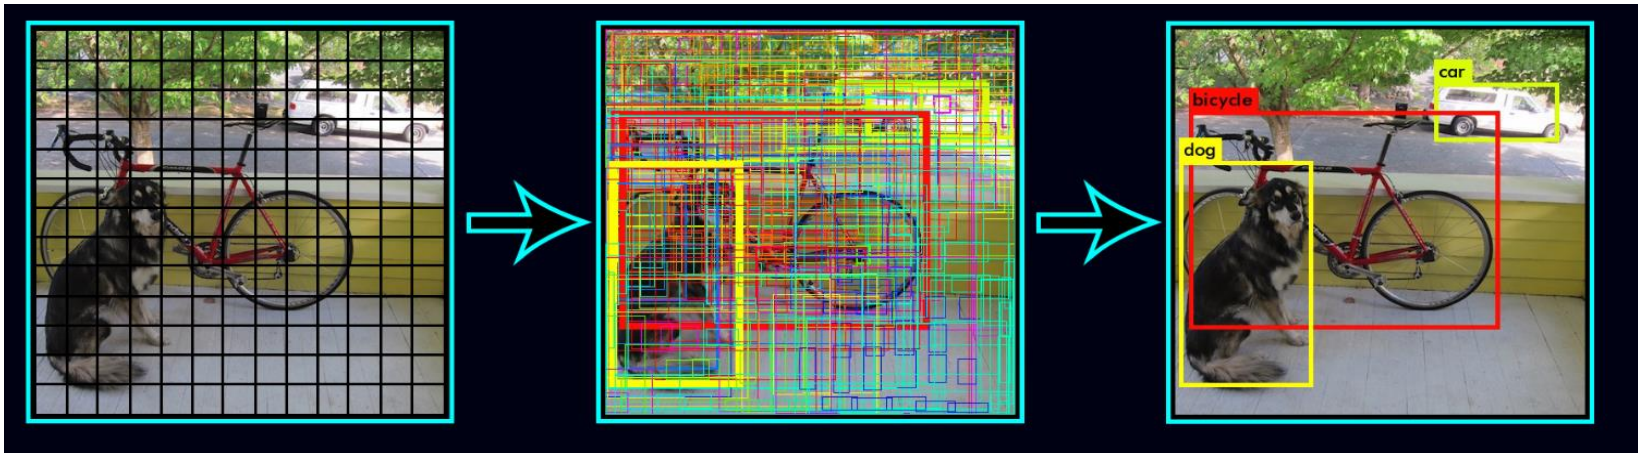

Source: https://pjreddie.com/darknet/yolov2/ (11.09.2018)

R-CNN

- Two-step proposal and refinement scheme
- Very large amount of proposals at high resolution needed

2. Symbol detection – methodology (contd.)
The deep watershed detector

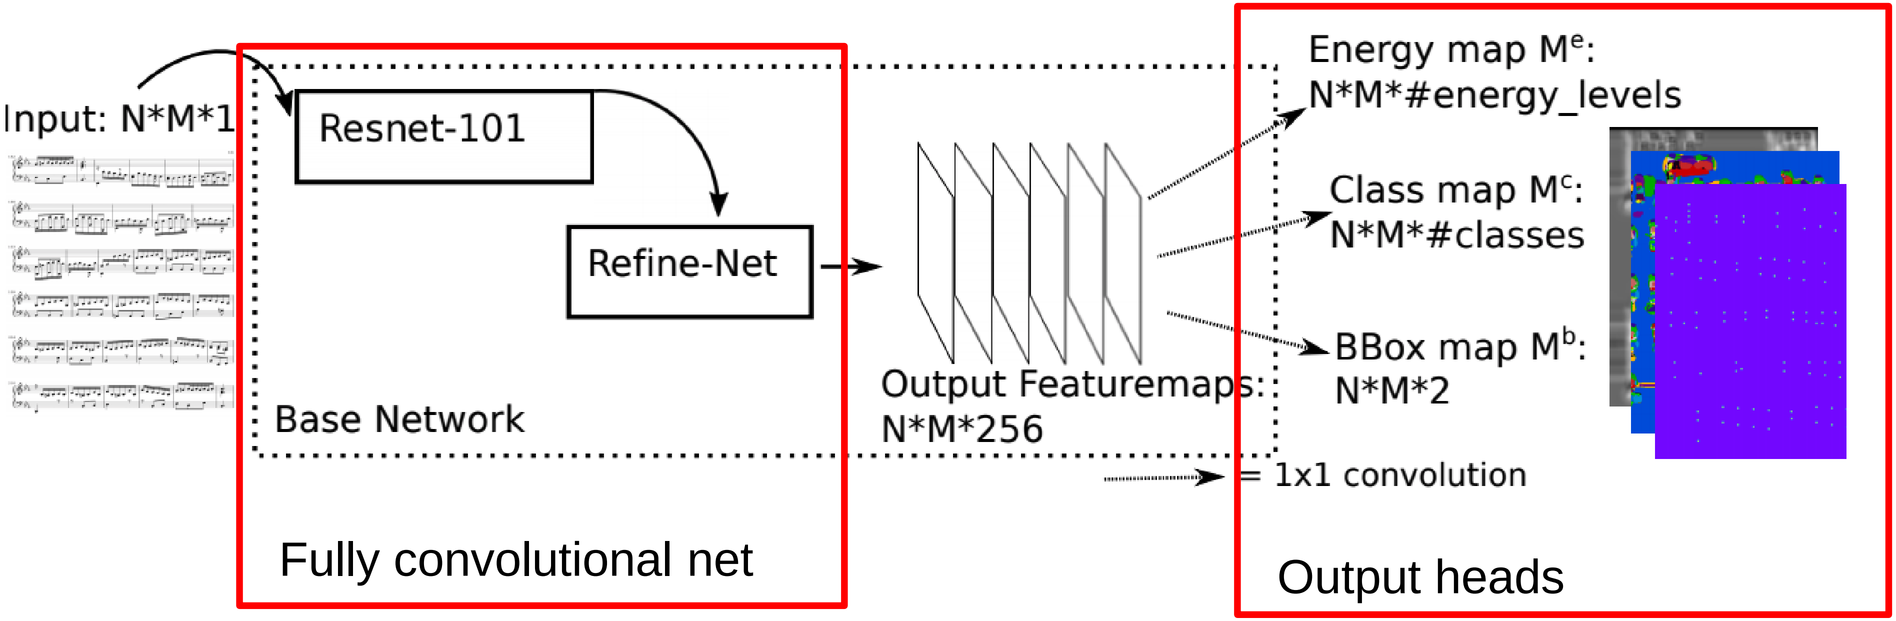

2. Symbol detection – methodology (contd.)

Output heads of the deep watershed detector

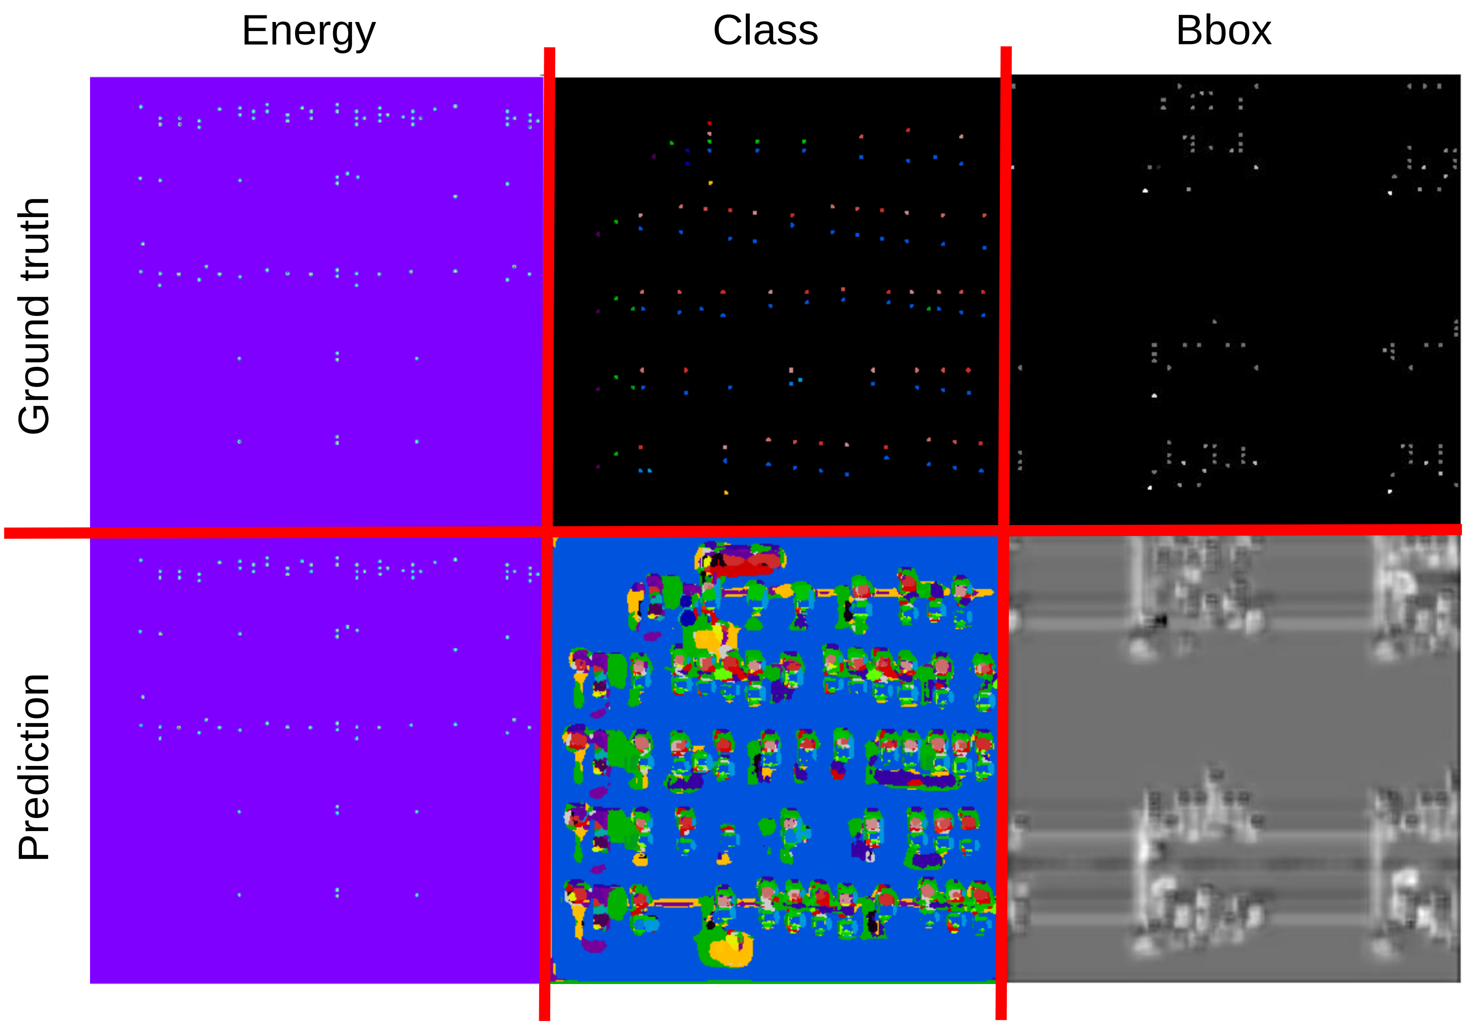

Zürcher Hochschule
für Angewandte Wissenschaften

2. Symbol detection – industrialization

Schweizerische Eidgenossenschaft
Confédération suisse
Confederazione Svizzera
Confederaziun svizra

Swiss Confederation

Innosuisse – Swiss Innovation Agency

Current results on class imbalance and robustness challenges

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

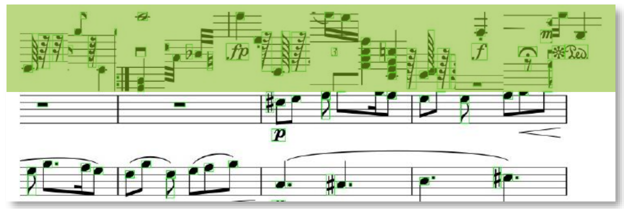

2. Put additional effort (and compute) into hyperparameter **tuning** and **longer training**
3. Trained also on scanned (more **real-worldish**) scores

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.

3. Face matching

Zürcher Hochschule
für Angewandte Wissenschaften

3. Face matching – challenges & solutions

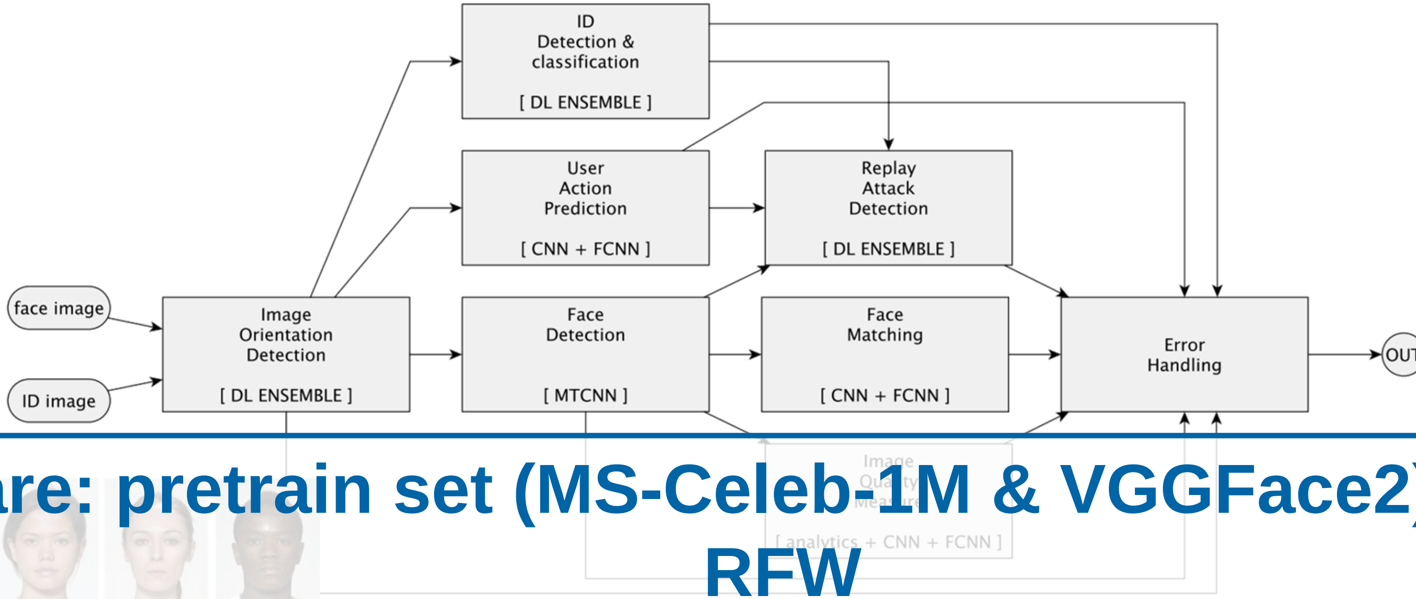

Beware: pretrain set (MS-Celeb-1M & VGGFace2) contains RFW

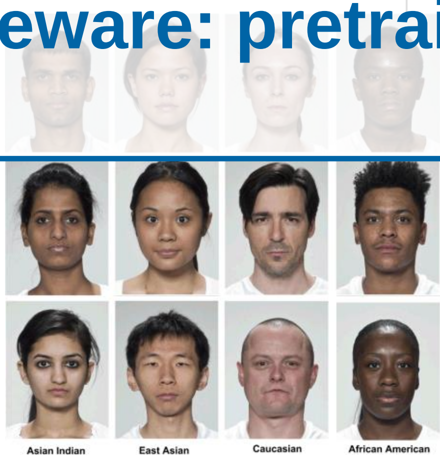

Test set: RFW	Industrial prototype	SE-ResNet-50 from scratch	SE-ResNet-50 pretrained	SE-ResNet-50 pretrained-128	SE-ResNet-50 pretrained-256
African	73.63	74.43	79.85	**80.20**	76.20
Asian	78.77	79.48	**83.44**	83.41	80.13
Caucasian	87.99	85.61	88.78	**89.76**	88.77
Indian	84.16	82.49	85.51	**85.97**	83.44
all	80.54	79.80	83.82	**84.35**	81.56

Stadelmann, Amirian, Arabaci, Arnold, Duivesteijn, Elezi, Geiger, Lörwald, Meier, Rombach & Tuggener (2018). *«Deep Learning in the Wild»*. ANNPR'2018.
Wang, Deng, Hu, Peng, Tao, & Huang (2018). *«Racial Faces in-the-Wild: Reducing Racial Bias by Deep Unsupervised Domain Adaptation»*. arXiv:1812.00194.
Hu, Shen & Sun (2018). *«Squeeze-and-Excitation Networks»*. CVPR'2018.

[!] DEEPIMPACT

Schweizerische Eidgenossenschaft
Confédération suisse
Confederazione Svizzera
Confederaziun svizra
Swiss Confederation
Innosuisse – Swiss Innovation Agency

4. Data-driven condition monitoring

Situation: Maintaining big (rotating) machinery is expensive, defect is more expensive

Goal: Schedule maintenance shortly before defect is expected, not merely regularly

Challenge: Develop an approach that adapts to each new machine automatically

Solution: Use machine learning approaches for **anomaly detection** to learn the normal state of each machine and deviations of it purely from observed sensor signals; the approach combines classic and industry-proven features with e.g. deep learning auto-encoders

vibration sensors

feature extraction

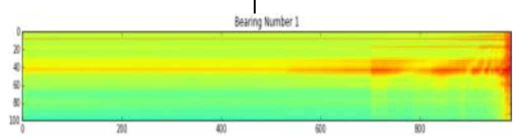

e.g., RNN autoencoder

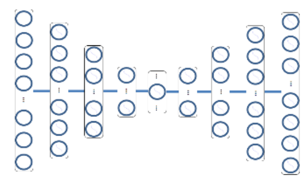

early detection of fault

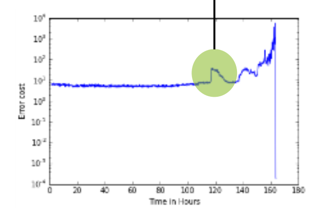

Stadelmann, Tolkachev, Sick, Stampfli & Dürr (2019): *«Beyond ImageNet–Deep Learning in Industrial Practice»*. In: Braschler et al. (Ed.), *«Appl. Dat. Sci.», Springer.*

4. Data-driven condition monitoring – results

Signal:

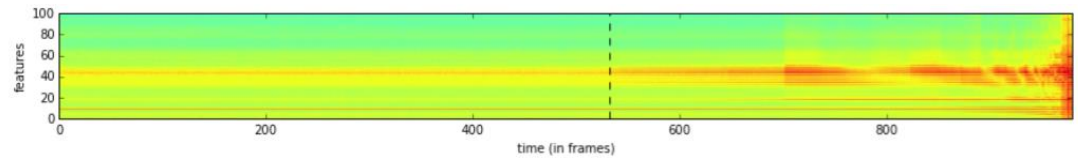

Shallow learning methods:

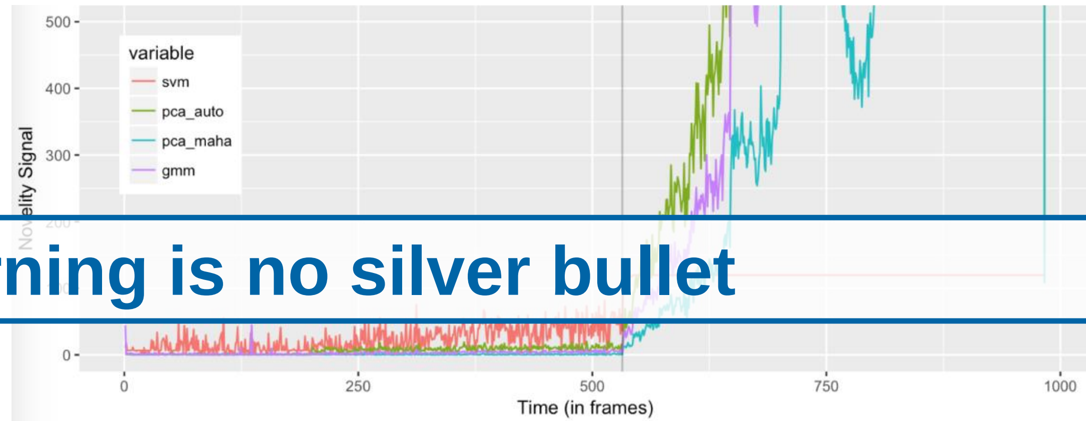

Deep learning is no silver bullet

Deep learning methods:

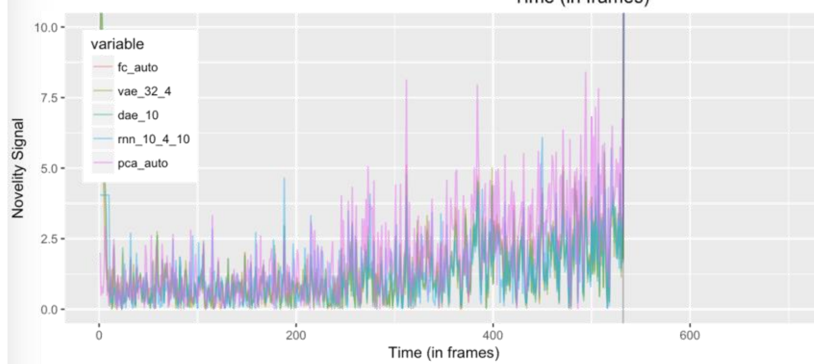

➔ DL and standard methods detect the defect time; DL show **less novelty** where there is **still no defect**

5. Speaker clustering

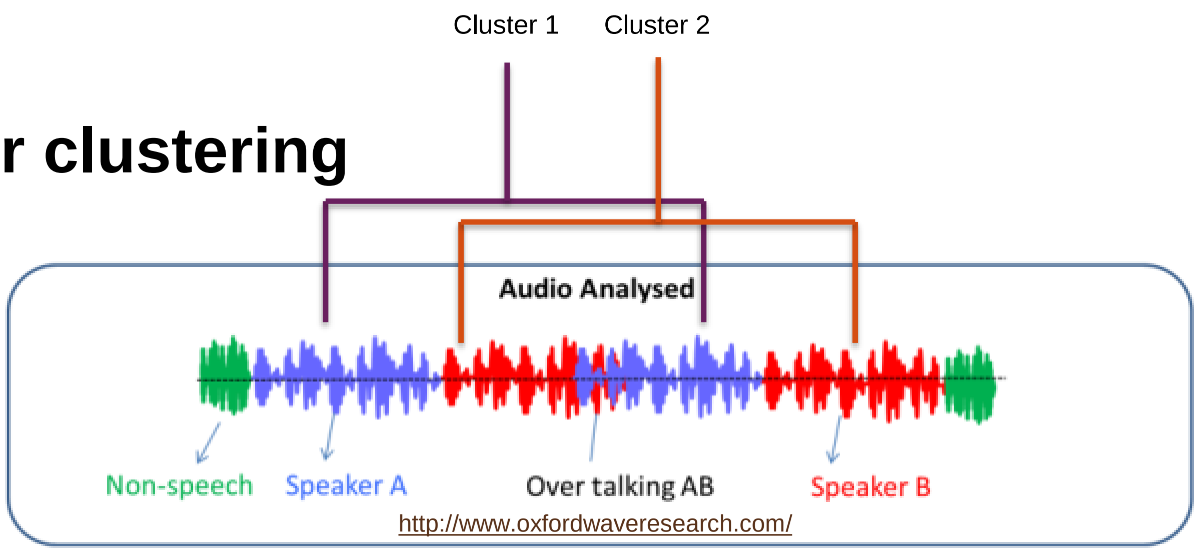

For the 630 training utterances, GMMs with 32 mixtures are built a priori, then an identification experiment is run for the 630 test utterances. It yields a satisfactory 0.5% closed set identification error.

[34]. Evaluations typically concentrate on data sets built from broadcast news/shows and meeting recordings, where diarization error rates ranging from 8% to 24% are reported [28][34][45]. These results are confirmed by more recent

The hypothesis of this paper is: the techniques originally developed for speaker verification and identification are not suitable for speaker clustering, taking into account the escalated difficulty of the latter task. However, the processing chain for speaker clustering is quite large – there are many potential areas for improvement. The question is: *where* should improvements be made to improve the *final* result?

The interpretation of our results has shown that it is the stage of modeling that bears the highest potential: the inclusion of temporal context information among feature vectors is what is crucially missing there. Furthermore, the inclusion

context vector. This corresponds to a syllable length of 130 ms and is found to best capture speaker specific sounds in informal listening experiments over a range of 32–496 ms (in intervals of 16 ms). Our context vector step is one orig-

Stadelmann & Freisleben (2009). *«Unfolding Speaker Clustering Potential: A Biomimetic Approach»*. ACMMM'2009.

5. Speaker clustering – exploiting time information

CNN (MLSP'16)

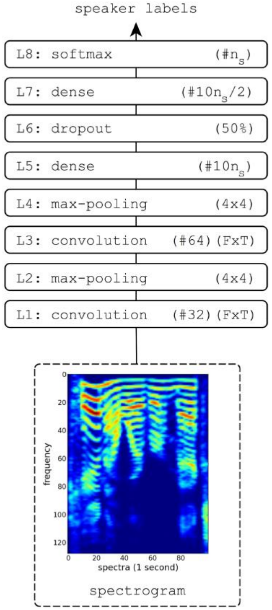

CNN & clustering-loss (MLSP'17)

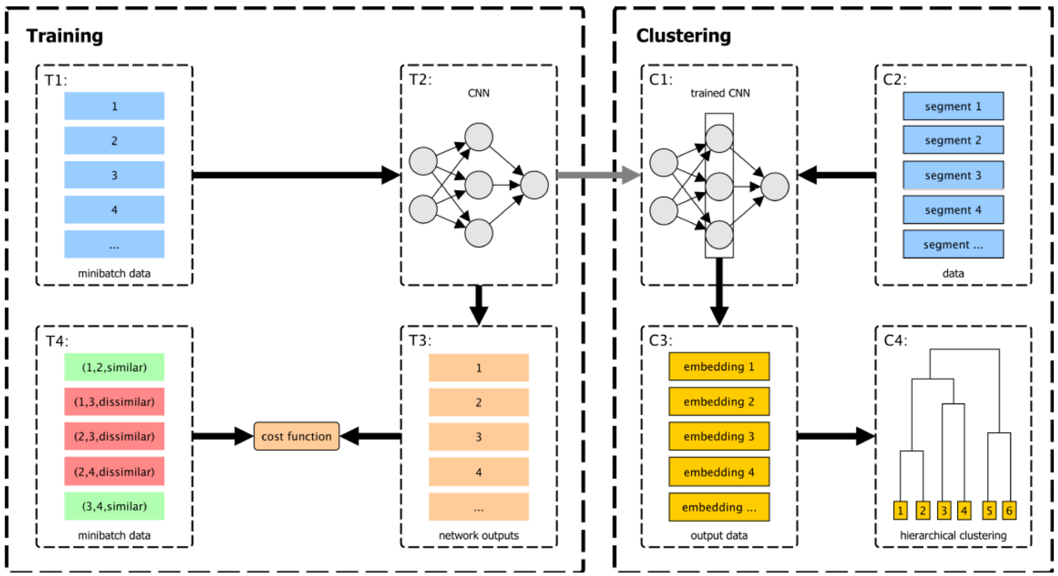

RNN & clustering-loss (ANNPR'18)

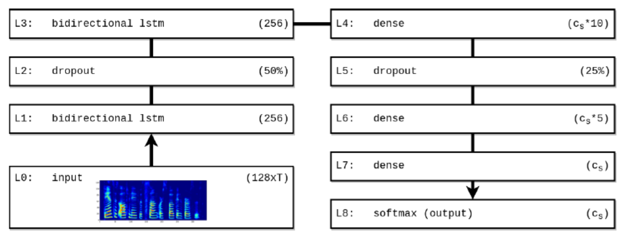

Method	MR	MR (legacy)
RNN /w PKLD	**2.19%** $\left(\frac{1.25\%+2.5\%+1.25\%+3.75\%}{4}\right)$	**4.38%** (average of 4 runs)
CNN /w PKLD [24]	-	**5%**
CNN /w cross entropy [23]	-	**5%**
ν-SVM [40]	6.25%	-
GMM/MFCC [40]	12.5%	-

Lukic, Vogt, Dürr & Stadelmann (2016). *«Speaker Identification and Clustering using Convolutional Neural Networks»*. MLSP'2016.
Lukic, Vogt, Dürr & Stadelmann (2017). *«Learning Embeddings for Speaker Clustering based on Voice Equality»*. MLSP'2017.
Stadelmann, Glinski-Haefeli, Gerber & Dürr (2018). *«Capturing Suprasegmental Features of a Voice with RNNs for Improved Speaker Clustering»*. ANNPR'2018.

5. Speaker clustering – methodology

Idea

- **Leverage** on recent success of **deep learning** in audio processing
- Use **RNN for** its known **sequence learning** capabilities
- **Extract** speaker **embeddings** for new utterance from trained RNN

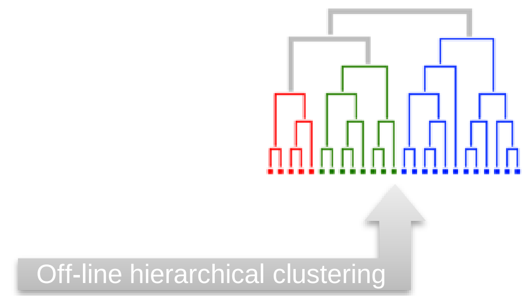

Output: Embedding (speaker-specific representation)

Model: Deep recurrent neural net

Input: Audio snippet

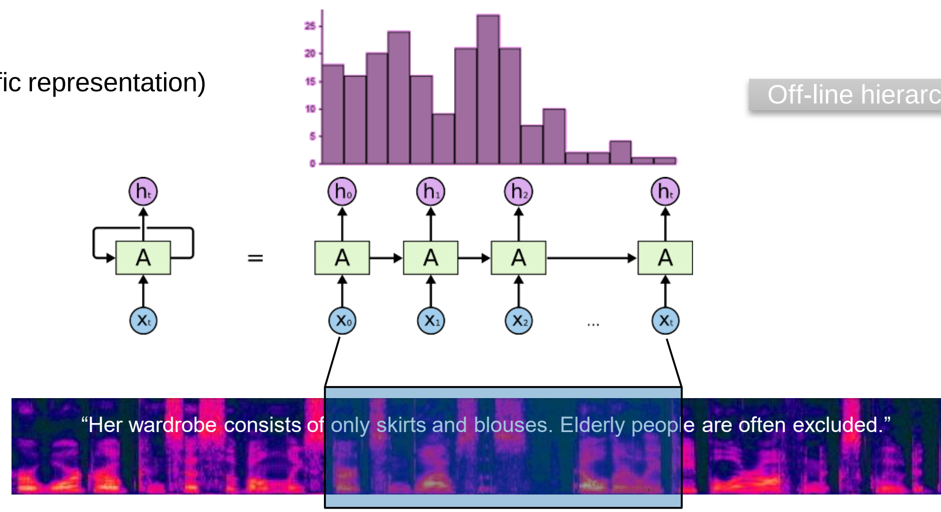

Challenges

- **RNNs** known to be **hard to train**
- Additionally: **no natural training target** → need surrogate task with hopefully helpful loss

5. Speaker clustering – methodology (contd.)

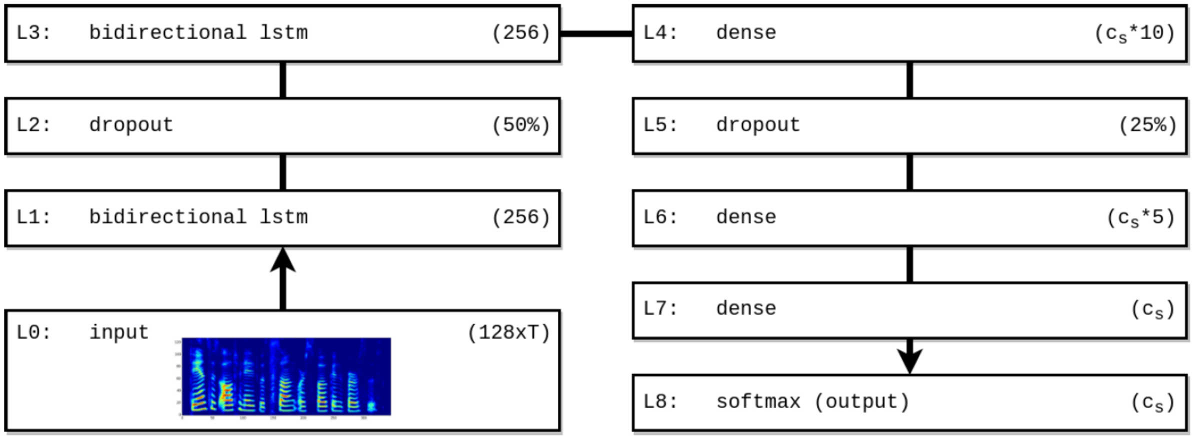

&

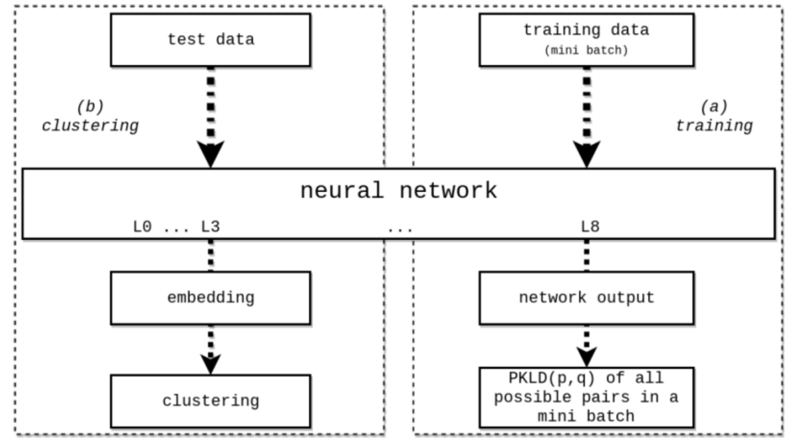

Learning target

- **Lxx** to output a **distribution** (c_S = number of speakers in training set) that is similar for samples of the same speaker, dissimilar for different speakers

Loss

- For all pairs (p, q) of distributions in a mini batch:
  - **Pairwise Kullback-Leibler** distance between **same-speaker** pairs:
  - **Hinge** loss (with hyperparameter $margin$) between **different-speaker** pairs:
- (final loss gets symmetrized)

$$\mathrm{KL}(\mathbf{p} \parallel \mathbf{q}) = \sum_i^{c_s} p_i \log \frac{p_i}{q_i}$$

$$\mathrm{HL}(\mathbf{p} \parallel \mathbf{q}) = \max(0, \mathrm{margin} - \mathrm{KL}(\mathbf{p} \parallel \mathbf{q}))$$

5. Speaker clustering – learnings & future work

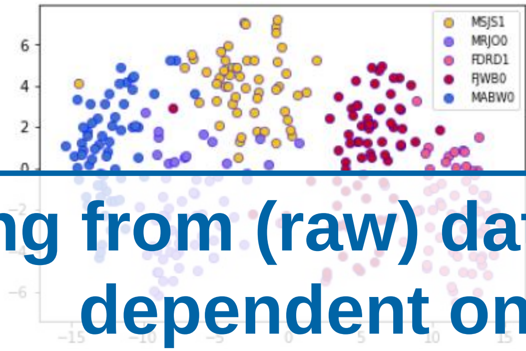

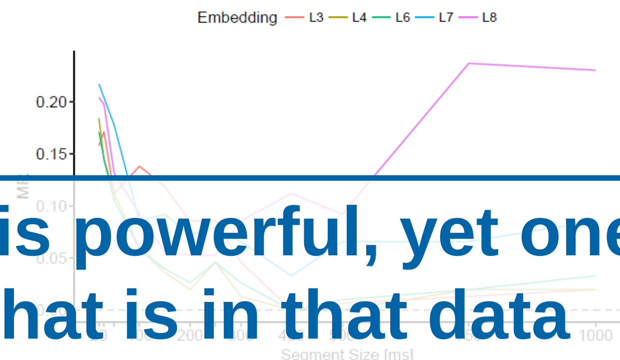

Learning from (raw) data is powerful, yet one is fully dependent on what is in that data

«Pure» voice modeling seems largely solved

- RNN **embeddings work well** (see t-SNE plot of single segments)
- RNN model robustly exhibits *the predicted* **«sweet spot» for** the used **time information**
- Speaker clustering on clean & reasonably long input works **an order of magnitude better** (*as predicted*)
- Additionally, using a smarter clustering algorithm on top of embeddings makes **clustering on TIMIT as good as identification** (see ICPR'18 paper on dominant sets)

Future work

- Make models robust on **real-worldish data** (noise and more speakers/segments)
- Exploit findings for robust reliable **speaker diarization**
- **Learn** embeddings and the clustering algorithm **end to end**

Hibraj, Vascon, Stadelmann & Pelillo (2018). *«Speaker Clustering Using Dominant Sets»*. ICPR'2018.
Meier, Elezi, Amirian, Dürr & Stadelmann (2018). *«Learning Neural Models for End-to-End Clustering»*. ANNPR'2018.

6. Lessons learned

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**

Prerequisite: stable data acquisition pipeline
Learning from (raw) data is powerful, yet one is fully dependent on what is in that data
Beware: pretrain set (MS-Celeb-1M & VGGFace2) contains RFW

Robustness is important

- **Training processes** can be tricky
  → give hints via a unique loss, proper preprocessing and pretraining

Sufficient condition: lots of tuning
Deep learning is no silver bullet

6. Lessons learned (contd.)

Interpretability is required and possible

- Helps the developer in «debugging», needed by the user to trust
  → visualizations of learned features, training process, learning curves etc. should be «always on»

negative X-ray

positive X-ray

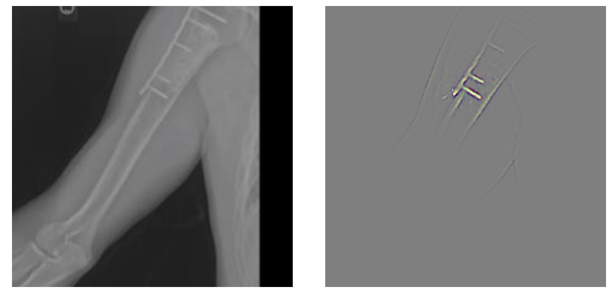

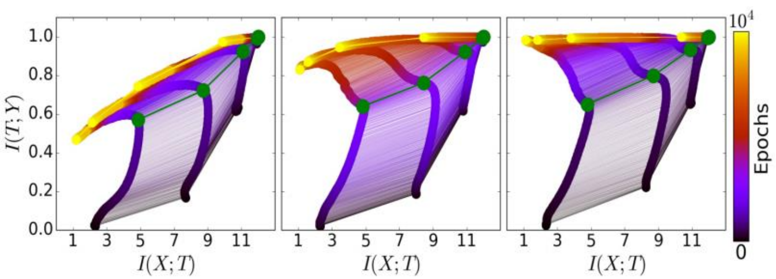

DNN training on the Information Plane

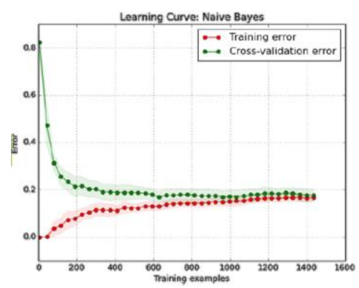

a learning curve

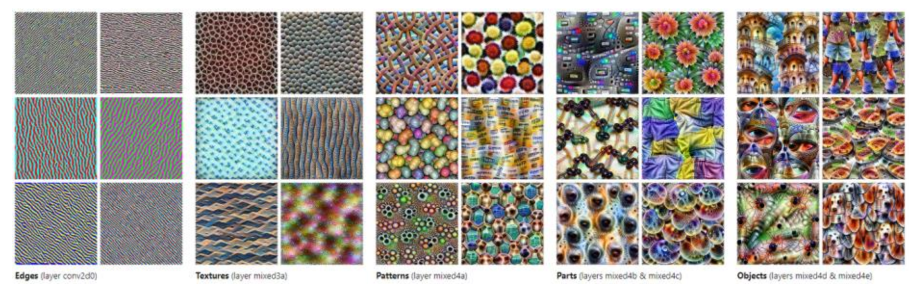

feature visualization

Stadelmann, Amirian, Arabaci, Arnold, Duivesteijn, Elezi, Geiger, Lörwald, Meier, Rombach & Tuggener (2018). *«Deep Learning in the Wild»*. ANNPR'2018.
Schwartz-Ziv & Tishby (2017). *«Opening the Black Box of Deep Neural Networks via Information»*.
https://distill.pub/2017/feature-visualization/, https://stanfordmlgroup.github.io/competitions/mura/

Zürcher Hochschule
für Angewandte Wissenschaften

6. Lessons learned – the greater good

The vision of Kai-Fu Lee, venture capitalist & scientist

- **AI** systems can take over **routine tasks**…
- …so that **humans** can follow their calling: **love** ("jobs of compassion")

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

Conclusions

- Deep learning **is applied** and deployed in «normal» businesses (non-AI, SME)
- **Data is key** (effort for acquisition and influence on results usually underestimated)
- DL **training** for new use cases **can be tricky** (→ thorough experimentation & tuning)
- AI's **challenge** is not so much **how we deal** with technology – but **with one another**

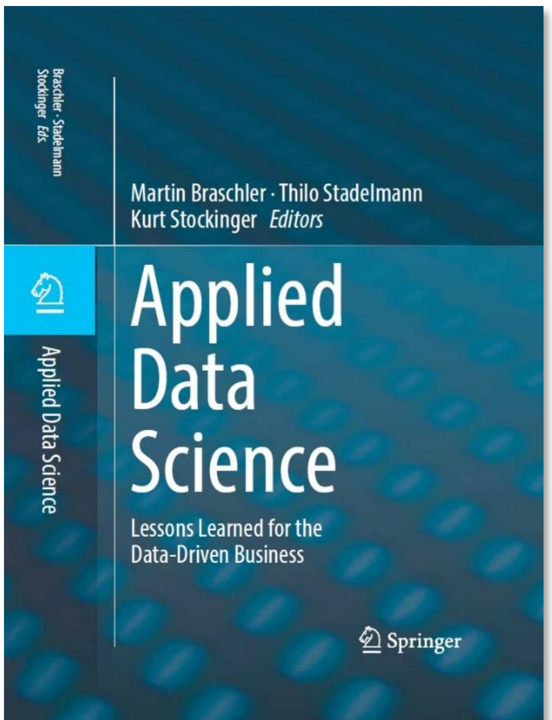

On me:

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

Further contacts:

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

➔ Happy to answer questions & requests.