Zürcher Hochschule
für Angewandte Wissenschaften

4 Years of Deep Learning Research @ ZHAW: an Information Engineering Perspective

InIT Meeting, Oct 17, 2018

T. Stadelmann, M. Amirian, I. Arabaci, M. Arnold, G. F. Duivesteijn, I. Elezi,
M. Geiger, S. Lörwald, B. B. Meier, K. Rombach & L. Tuggener

How time flies…
First Inspiration @ Zurich ML Meetup #1, Feb 25, 2014
First presentation of DL activities @ SP IE, Nov 14, 2014

How time flies…

First Inspiration @ Zurich ML Meetup #1, Feb 25, 2014
First presentation of DL activities @ SP IE, Nov 14, 2014

Thema «Deep Learning»

Aktuell «heissestes Thema» im Machine Learning
- (Aktueller) Kick-Off: Bengio et al., «A fast learning algorithm for deep belief nets», 2006
- Seitdem: Viele Pattern Recognition Benchmarks (teils um Grössenordnungen) durch DL-Ansätze verbessert (z.B. Schmidhuber, «Deep Learning in Neural Networks: An Overview», 2014)

Technisches
- Neuronales Netz mit vielen Schichten («deep»), einzeln vortrainiert z.B. als Autoencoder, ganzheitlich feinabgestimmt durch Backpropagation
- Spezialität: «Unsupervised Feature Learning» → Verfahren lernt selbständig Hierarchien «guter» (d.h. ähnlich wie der Mensch) Repräsentationen der Daten
- Anwendung: z.Zt. Pattern Recognition auf grossen Datenmengen (z.B. Bilder, Text, Audio)

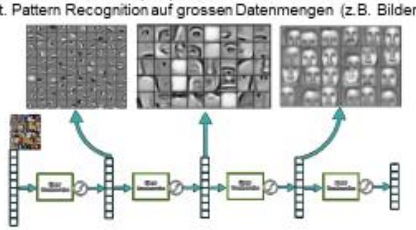

Impact? Markt? Research?

Strategische Bedeutung
- Vor 2 Jahren etwa <10 Forschungsgruppen an Top-Unis (u.a. IDSIA in der Schweiz)
- Januar 2014: Google kauft Fa. DeepMind für 500 Mio. $ (Gründung u.a. von Postdocs IDSIA)
- Aktuell:
  - Wissenschaftlich an der Grenze zwischen Forschung und Anwendung → stark abhängig von Anwendungsfall
  - Bücher, Libraries, Firmen entstehen erst
- Prognose: in 2-4 Jahren ein Tool am Markt wie «SVM»

Deep Learning an der ZHAW
- Angeschaffte Hardware: 2 Power-Workstations mit leistungsstarken GPUs «ready» (InIT & IDP)
- Deep Learning Journals Club: 12 Personen lesen gemeinsam UFLDL Tutorial aus Stanford
- Mailingliste: deeplearning@downbirn.zhaw.ch (für den Journals Club)
- DL Literatur-Archiv: https://www.dropbox.com/sh/duw20to4ugec73t/AADCPB9xoHCMlUmK1y8tueG8a?dl=0
- Verschiedene Ideen für interne Anträge (dueo am IDP; stdm mit baud; ciel mit uzdi)
- Erste KTI Anträge entstehen (acke mit stdm)

→ Für stdm/musy/stmf, den Schwerpunkt und das Datalab ein Fokusthema für die nächsten Jahre

Zürcher Hochschule
für Angewandte Wissenschaften

How time flies…
First Inspiration @ Zurich ML Meetup #1, Feb 25, 2014
First presentation of DL activities @ SP IE, Nov 14, 2014

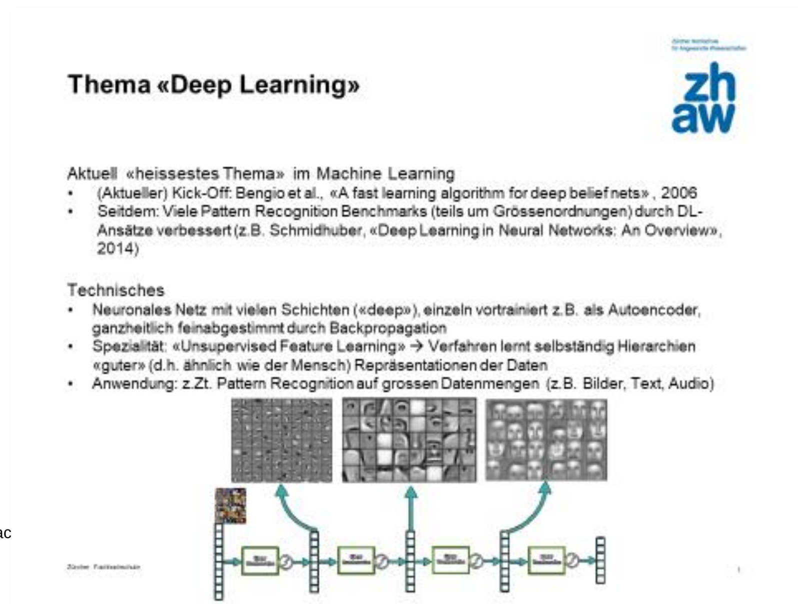

How time flies…

First Inspiration @ Zurich ML Meetup #1, Feb 25, 2014
First presentation of DL activities @ SP IE, Nov 14, 2014

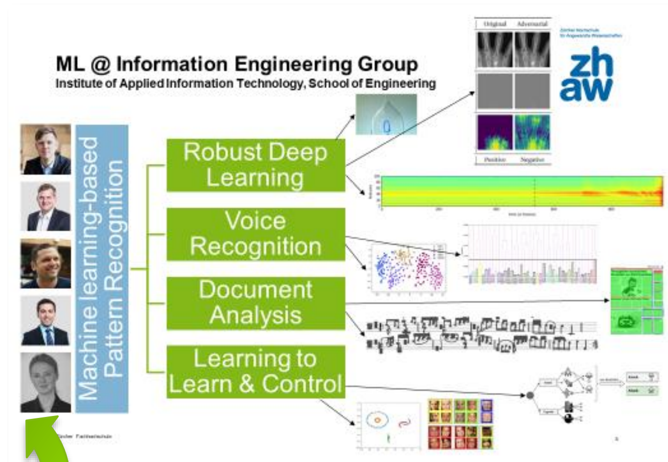

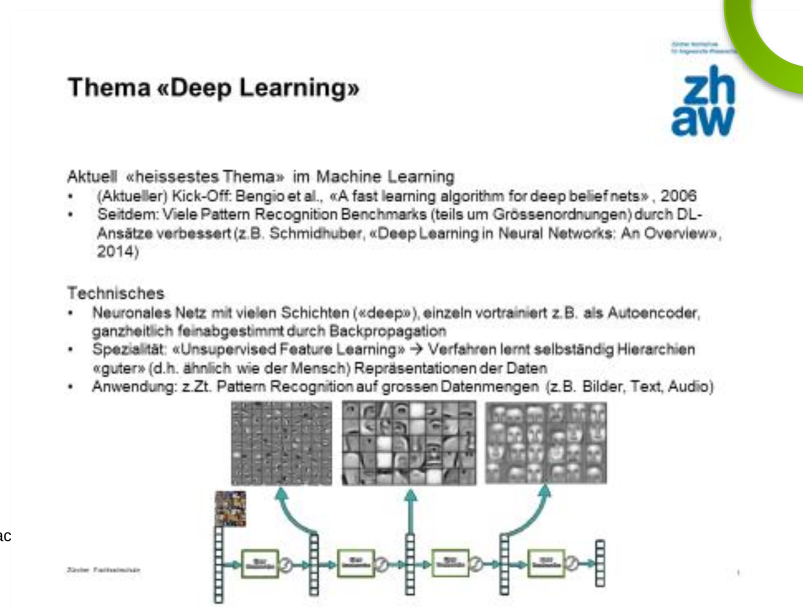

Research in the wild and in the lab

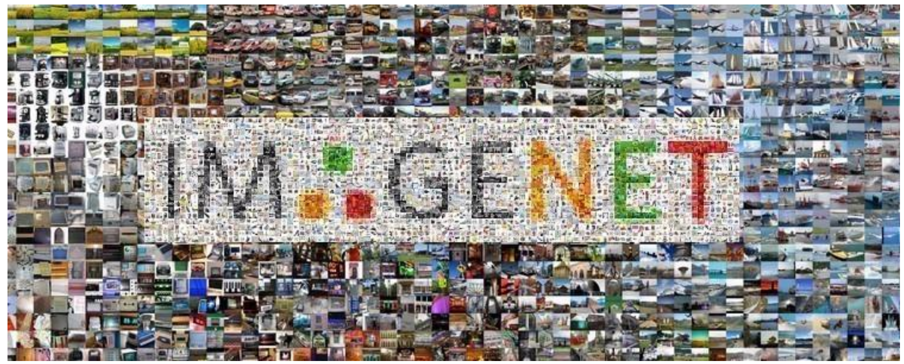

Research in the wild and in the lab

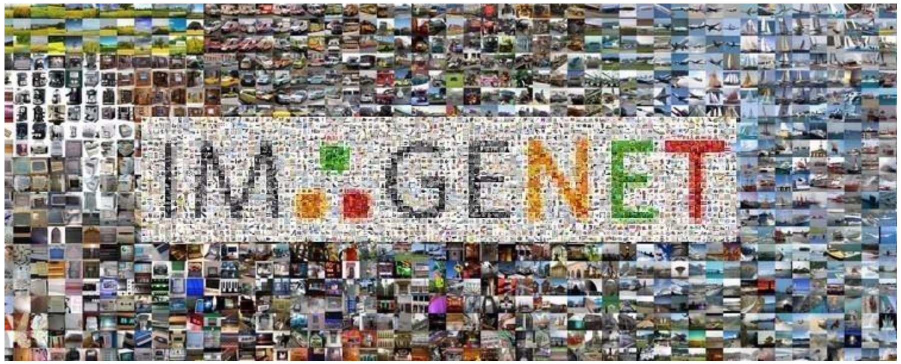

Motivated by general progress

- Given *known environment (*learning target, data, evaluation metric)

➔ Goal: *fundamental advance* in method

Research in the wild and in the lab

Motivated by application
- Facing *unclear*/unprecedented *learning target & data quality / quantity* issues

➔ Goal: new *product & advance* in method

Motivated by general progress
- Given *known environment* (learning target, data, evaluation metric)

➔ Goal: *fundamental advance* in method

Research in the wild and in the lab

Motivated by application

- Facing *unclear*/unprecedented *learning target & data quality / quantity* issues

➔ Goal: new *product & advance* in method

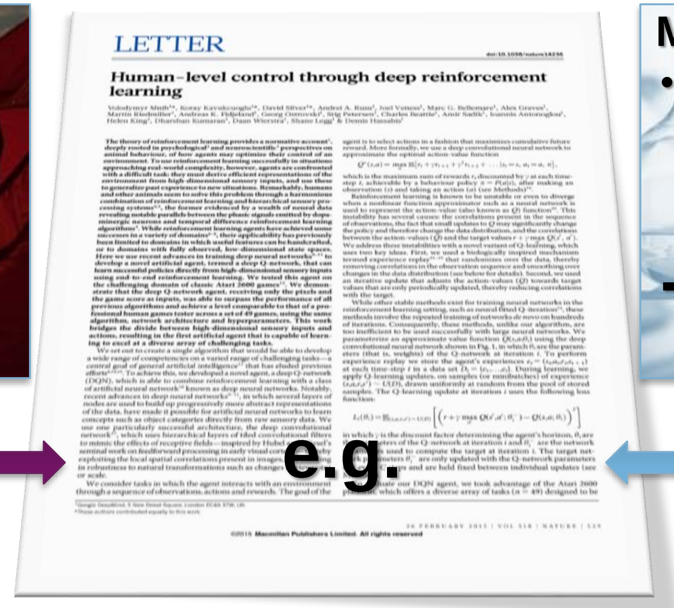
LETTER

doi:10.1038/nature14236

Human-level control through deep reinforcement learning

e.g.

Motivated by general progress

- Given *known environment* (learning target, data, evaluation metric)

➔ Goal: *fundamental advance* in method

Roadmap

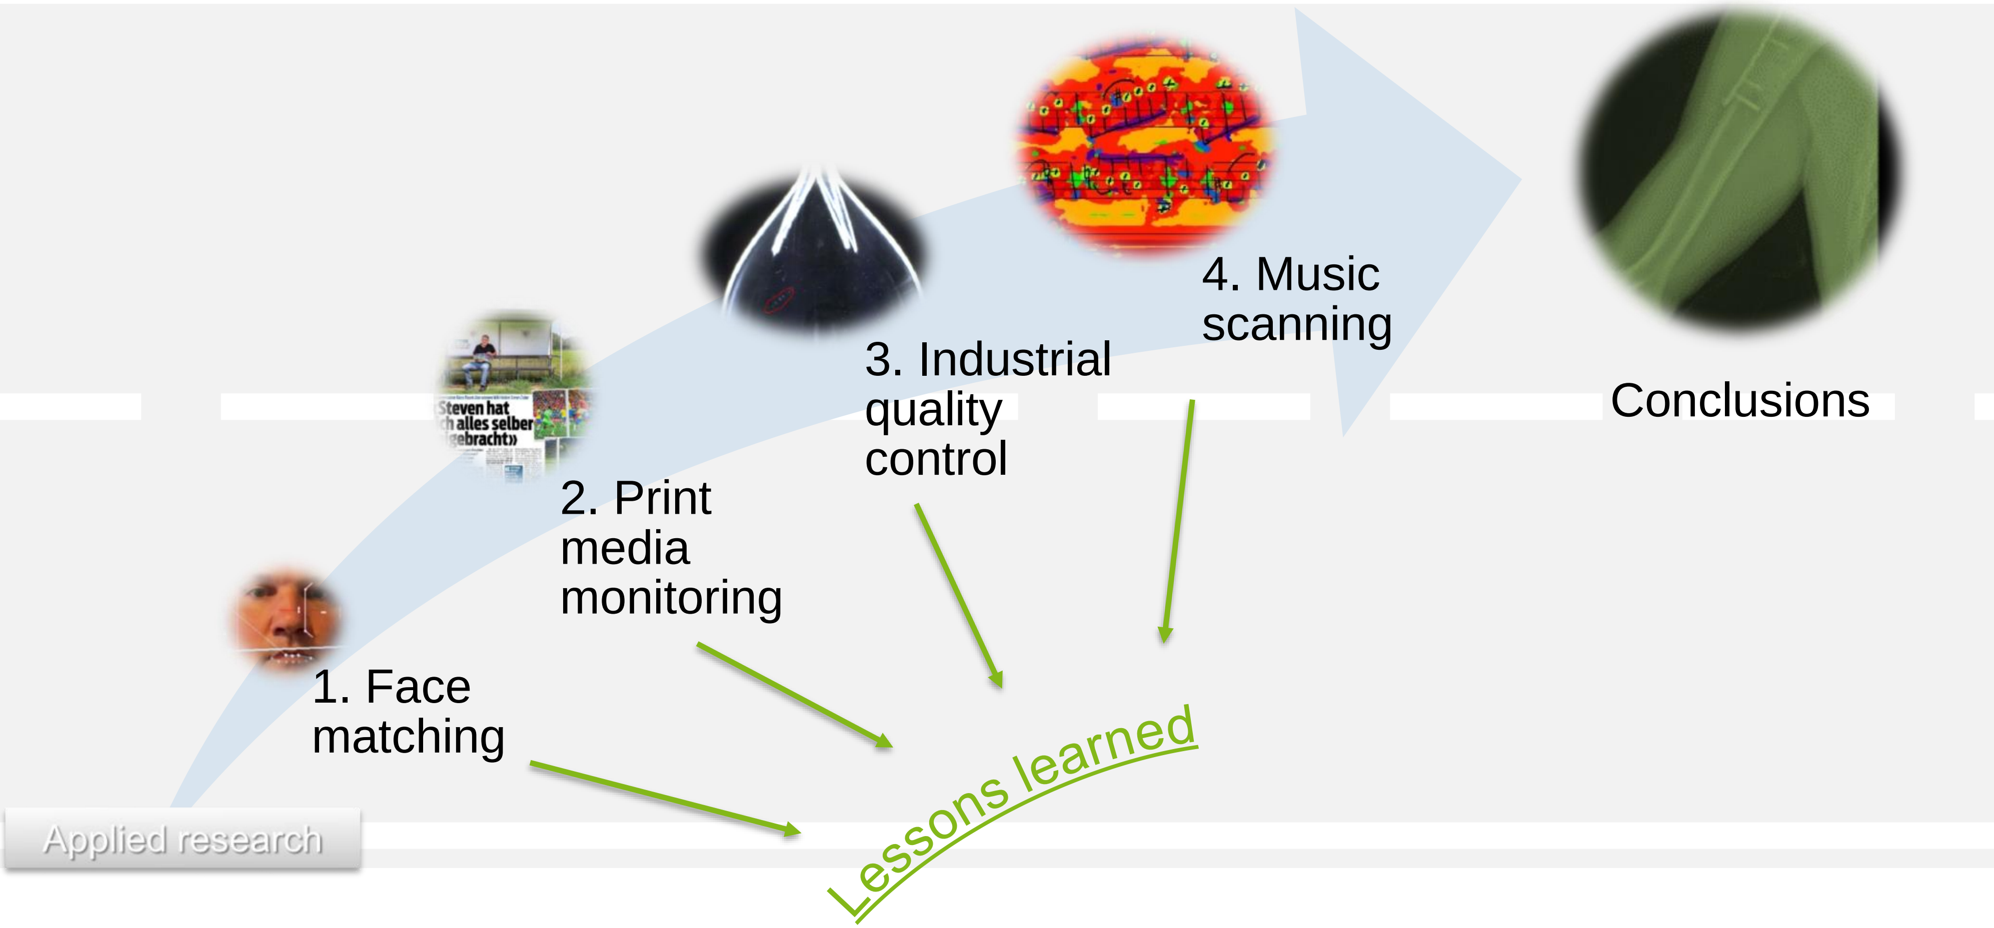
1. Face matching
2. Print media monitoring
3. Industrial quality control
4. Music scanning
Conclusions
Applied research
Lessons learned

1. Face matching

1. Face matching – challenges & solutions

Zürcher Hochschule
für Angewandte Wissenschaften

[!] DEEPIMPACT

Schweizerische Eidgenossenschaft
Confédération suisse
Confederazione Svizzera
Confederaziun svizra

Swiss Confederation

Innosuisse – Swiss Innovation Agency

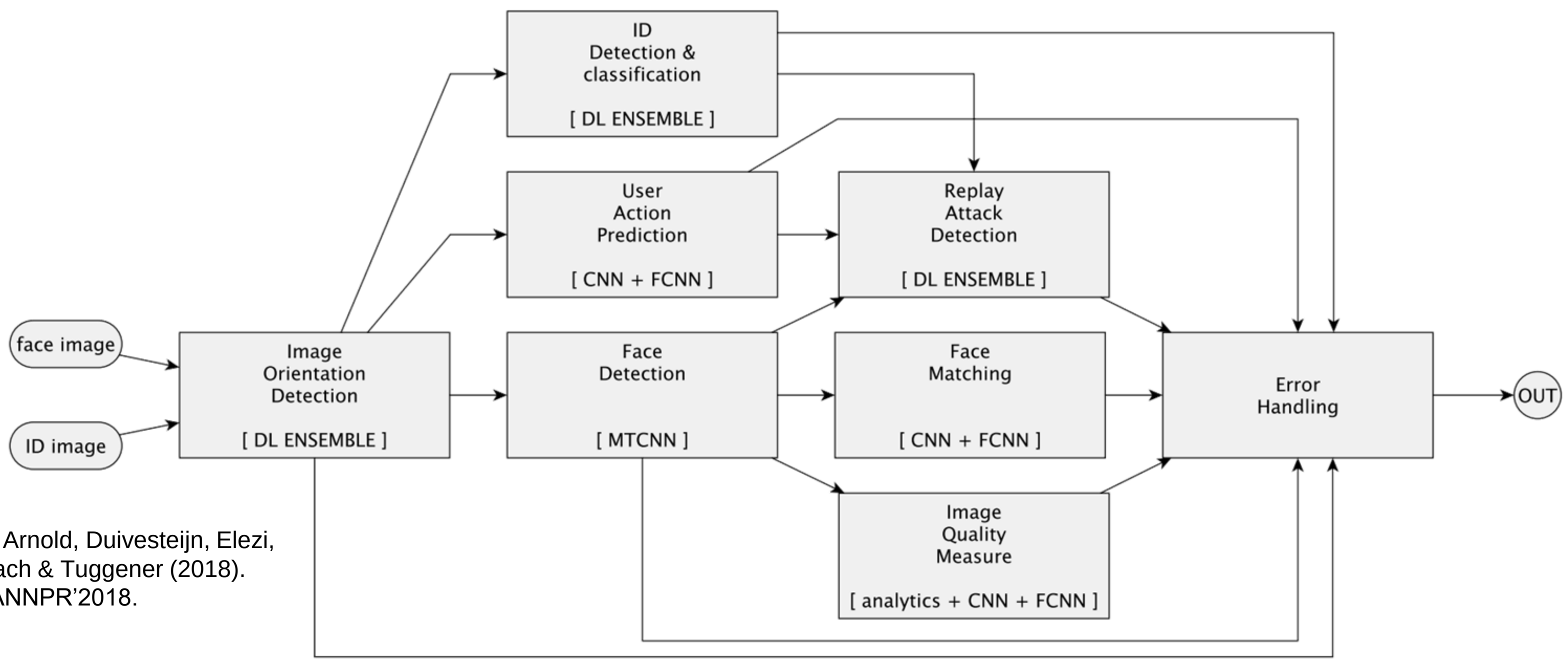

Stadelmann, Amirian, Arabaci, Arnold, Duivesteijn, Elezi, Geiger, Lörwald, Meier, Rombach & Tuggener (2018). *«Deep Learning in the Wild»*. ANNPR'2018.

1. Face matching – challenges & solutions



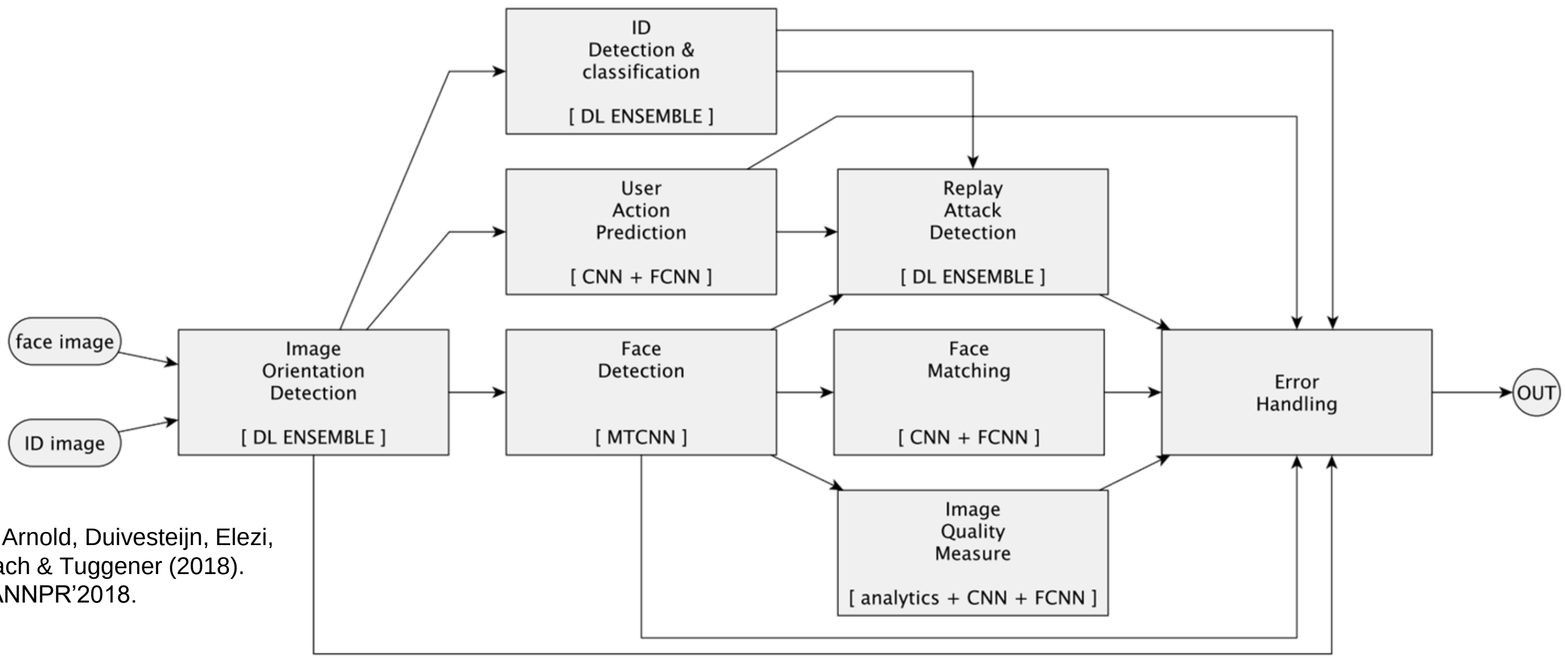

Stadelmann, Amirian, Arabaci, Arnold, Duivesteijn, Elezi, Geiger, Lörwald, Meier, Rombach & Tuggener (2018). *«Deep Learning in the Wild»*. ANNPR'2018.

1. Face matching – challenges & solutions

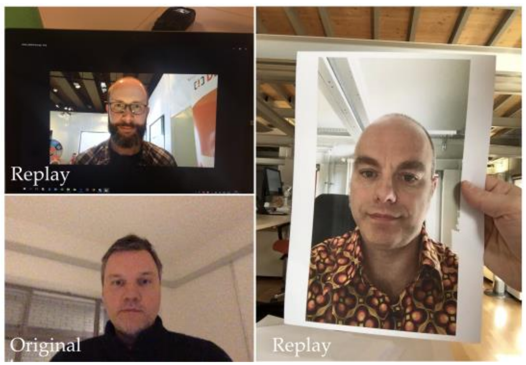

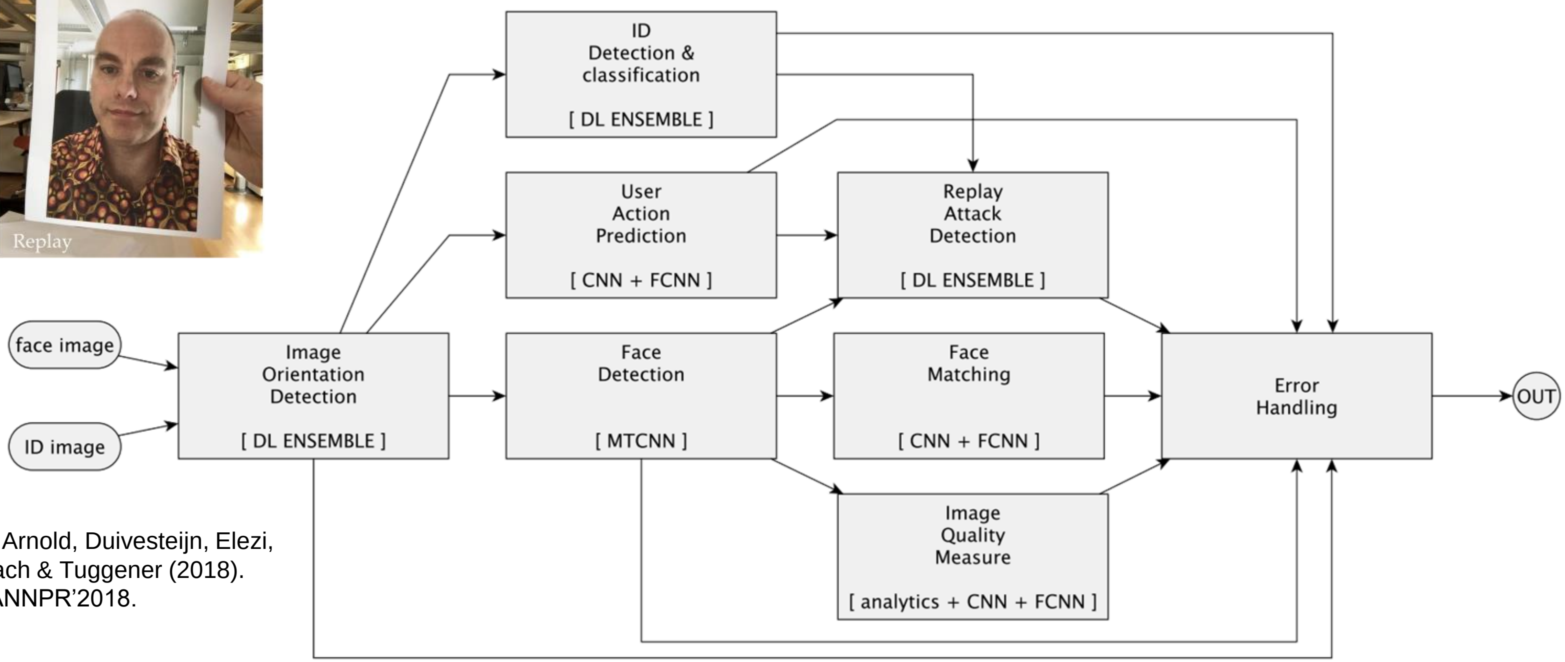

Stadelmann, Amirian, Arabaci, Arnold, Duivesteijn, Elezi, Geiger, Lörwald, Meier, Rombach & Tuggener (2018). *«Deep Learning in the Wild»*. ANNPR'2018.

[!] DEEPIMPACT

Schweizerische Eidgenossenschaft
Confédération suisse
Confederazione Svizzera
Confederaziun svizra

Swiss Confederation

Innosuisse – Swiss Innovation Agency

Lessons learned 1/4

,

Deployment

- Might involve the buildup of up to dozens of **other machine learning** models to flank the original core part.

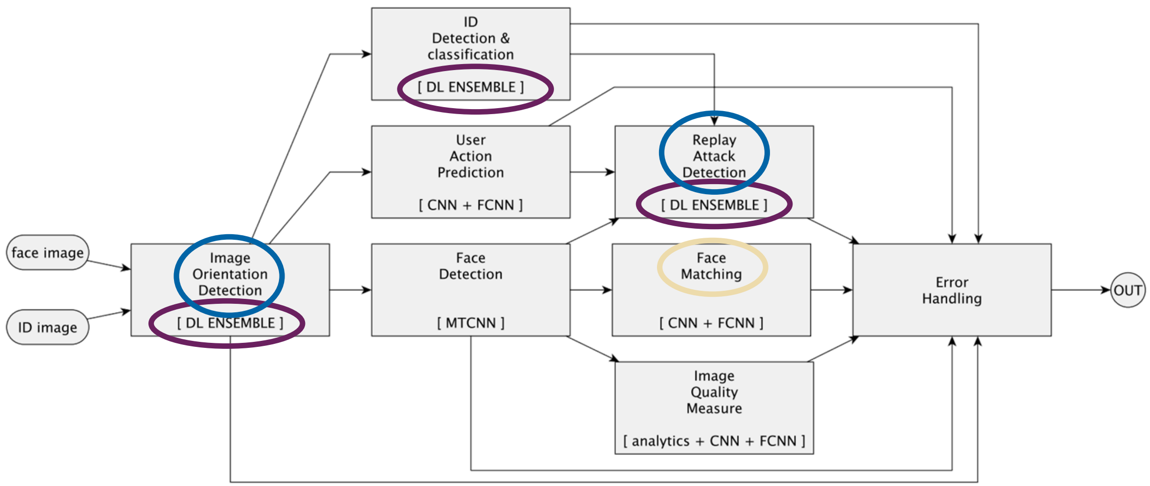

- **Specialized models** for identifiable sub-problems increase the accuracy in production systems over all-in-one solutions, and **ensembles** of experts help where no single method reaches adequate performance.

Zürcher Hochschule
für Angewandte Wissenschaften

2. Print media monitoring

Task

International.

Spionage für den Erzfeind Iran

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

Asylbewerber können bleiben

Nordkoreanischer Diktator zu Besuch in Peking

Vermögen beschlagnahmt

Klage von Le Pen abgelehnt

Challenge

Sein Juniorentrainer Mano Pavesi über unseren WM-Helden Steven Zuber

«Steven hat sich alles selber beigebracht»

Unsere Nati ist gegen Serbien unter Druck

Verlieren verboten!

Transfer Ticker

Liverpool will Yann Sommer!

Nuisance

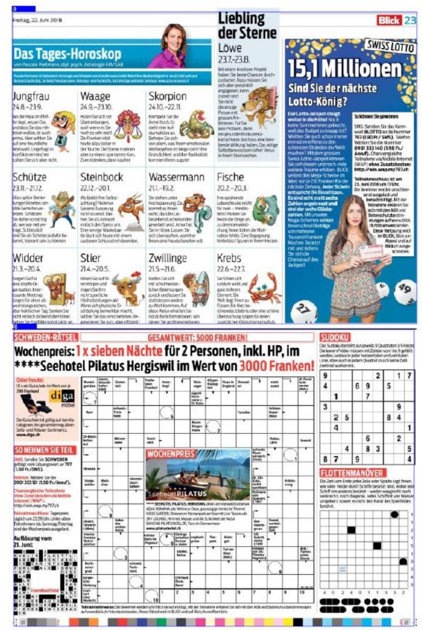

Das Tages-Horoskop

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

ARGUS DATA INSIGHTS
WISSEN ZUM ERFOLG

Schweizerische Eidgenossenschaft
Confédération suisse
Confederazione Svizzera
Confederaziun svizra

Swiss Confederation

Innosuisse – Swiss Innovation Agency

2. Print media monitoring – ML solution

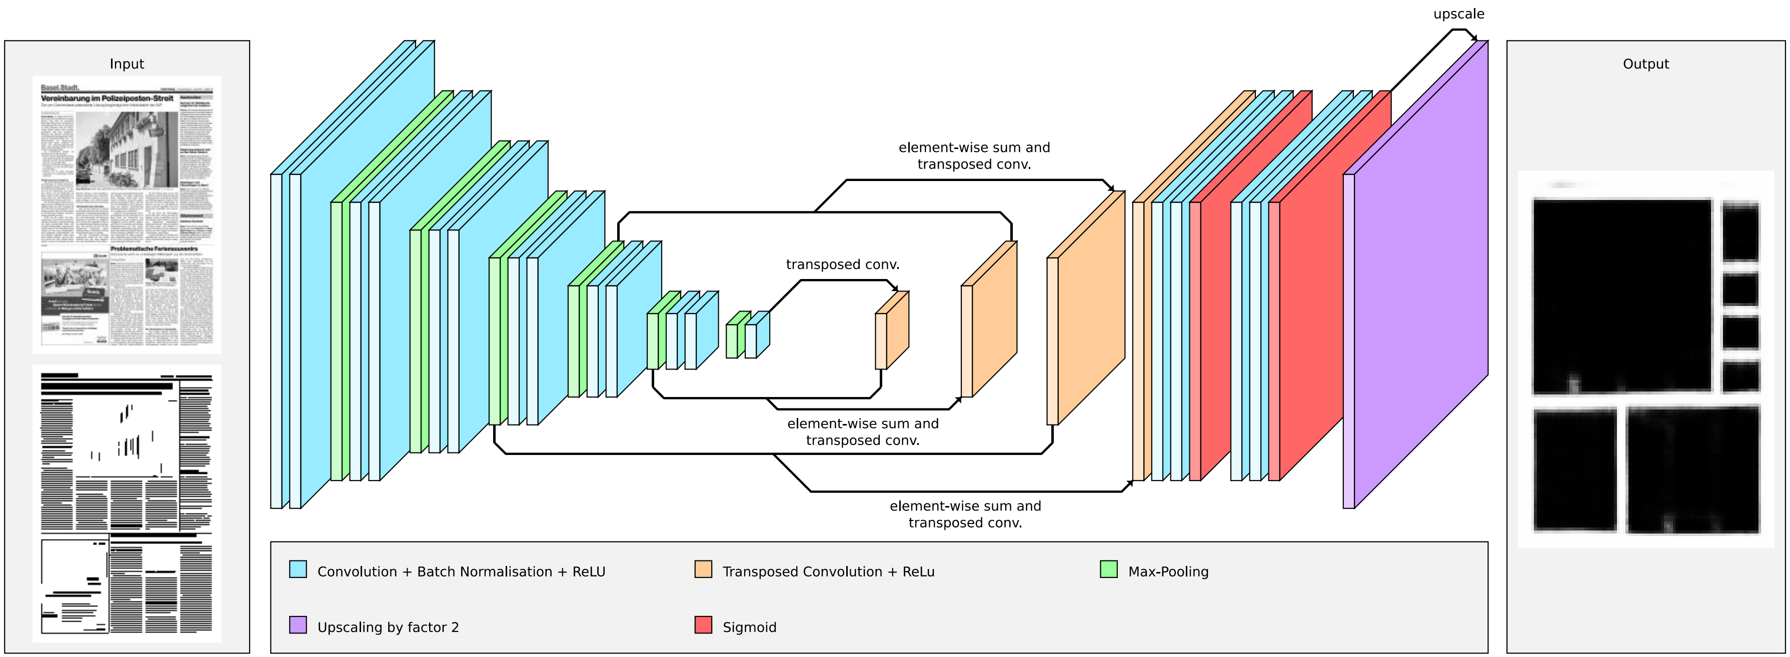

Meier, Stadelmann, Stampfli, Arnold & Cieliebak (2017). «*Fully Convolutional Neural Networks for Newspaper Article Segmentation*». ICDAR'2017.
Stadelmann, Tolkachev, Sick, Stampfli & Dürr (2018). «*Beyond ImageNet - Deep Learning in Industrial Practice*». In: Braschler et al., «Applied Data Science»,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

2. Print media monitoring – deployment

Zürcher Hochschule für Angewandte Wissenschaften

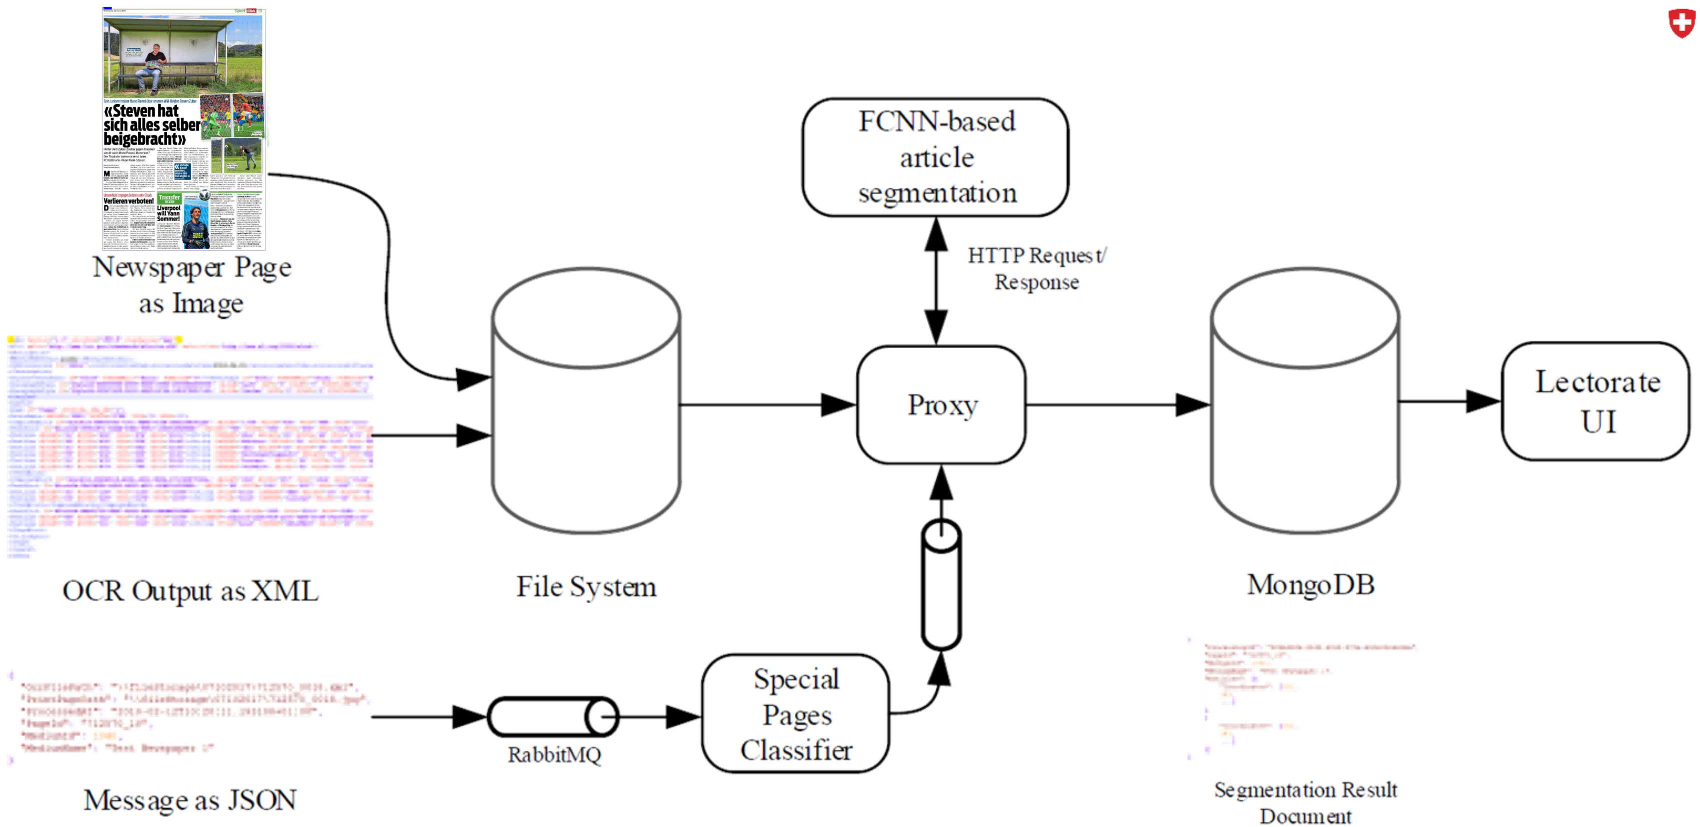

Stadelmann, Amirian, Arabaci, Arnold, Duivesteijn, Elezi, Geiger, Lörwald, Meier, Rombach & Tuggener (2018). *«Deep Learning in the Wild»*. ANNPR'2018.

Lessons learned 2/4

 , …

Deployment

- Should include **continuous learning**

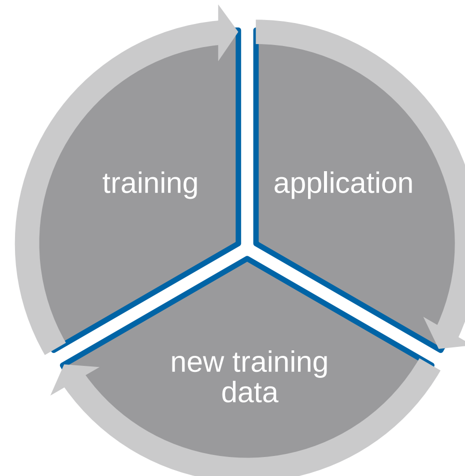

- Needs to take care of **processing speed** / efficiency

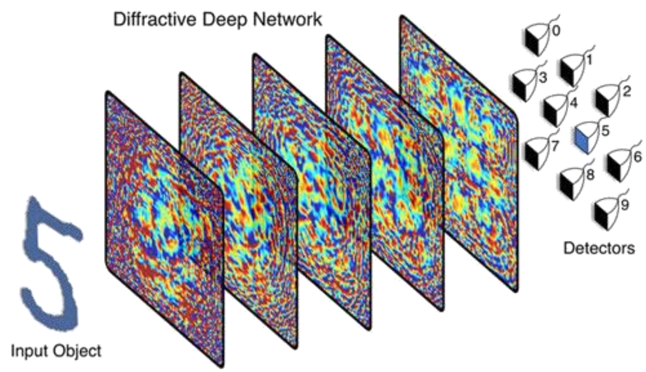

Symbolic image: a CNN in (optical) hardware (Lin et al., 2018).

Lin, Rivenson, Yardimci, Veli, Luo, Jarrahi & Oczan (2018). *«All-optical machine learning using diffractive deep neural networks»*. Science, 26. Jul 2018.

3. Industrial quality control

Task

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

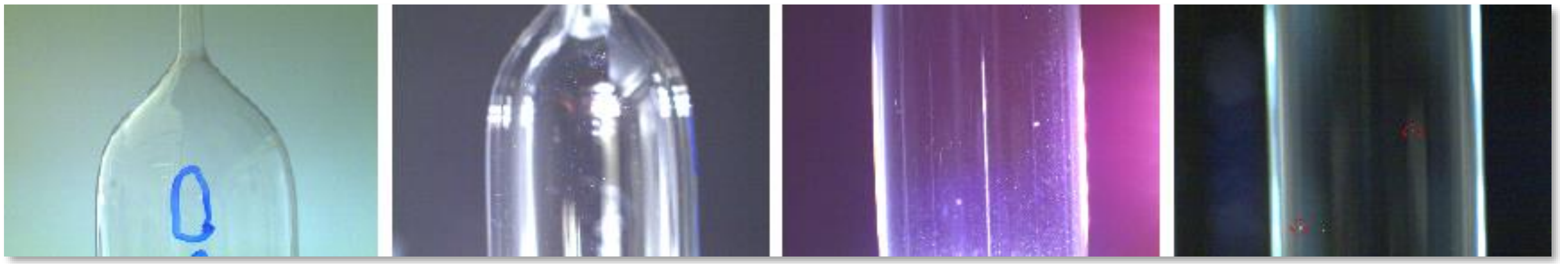

Challenges

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

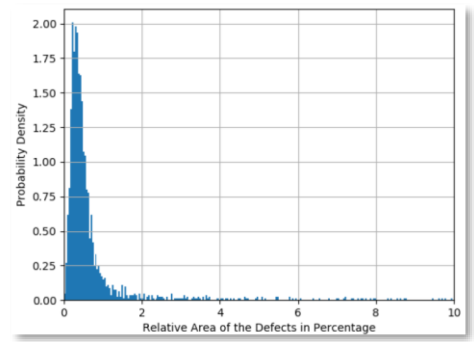

3. Industrial quality control – solutions
(Work in progress)

Ingredients
- Weighted loss
- Defect cropping
- Secret sauce

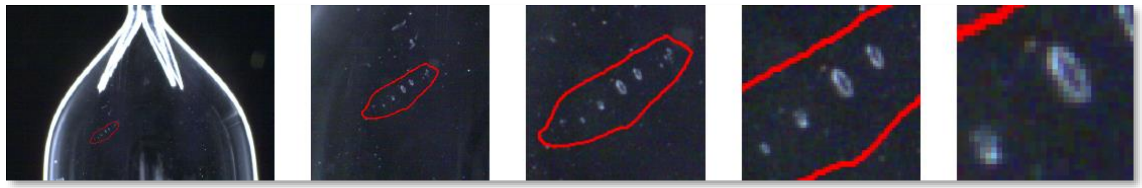

BW-TEC
Schweizerische Eidgenossenschaft
Confédération suisse
Confederazione Svizzera
Confederaziun svizra
Swiss Confederation
Innosuisse – Swiss Innovation Agency

Preliminary results

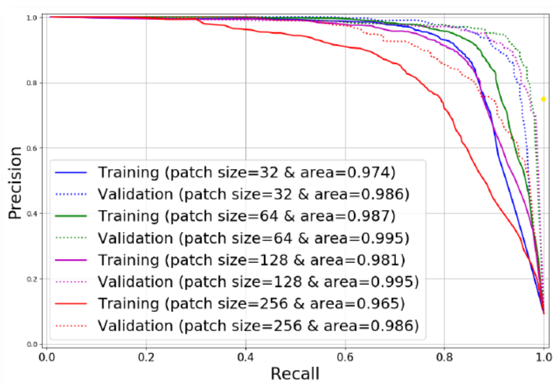

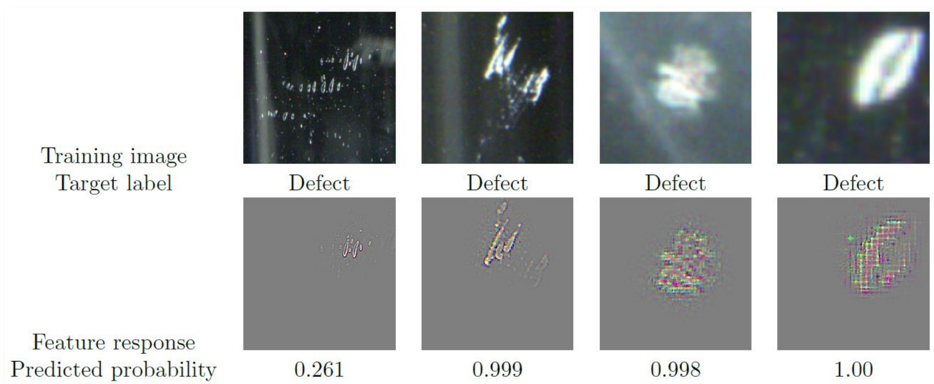

Lessons learned 3/4

, , 

Data

- Acquisition usually **needs much more time** than expected, yet is the basis for all subsequent success
- Class **imbalance** & **covariate shift** are usual

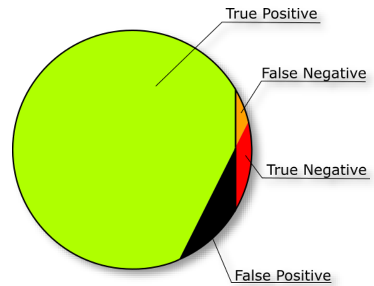

Lessons learned 3/4

, , 

Data

- Acquisition usually **needs much more time** than expected, yet is the basis for all subsequent success
- Class **imbalance** & **covariate shift** are usual

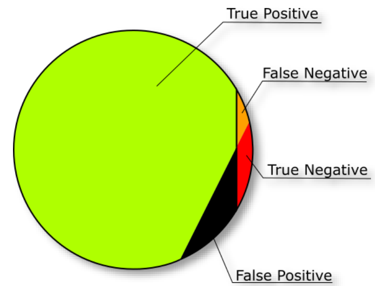

Understanding

- **What has been learned** and **how decisions emerge** help both the user and the developer of neural networks to build trust and improve quality

 

- **Operators and business owners** need a basic understanding of used methods to produce usable ground truth and provide relevant subject matter expertise

Lessons learned 3/4 (contd.)

, , , 

Simple baselines

- Do a good job in **determin**ing the **feasibility** as well as the **potential** of the task at hand when final datasets or novel methods are not yet seen

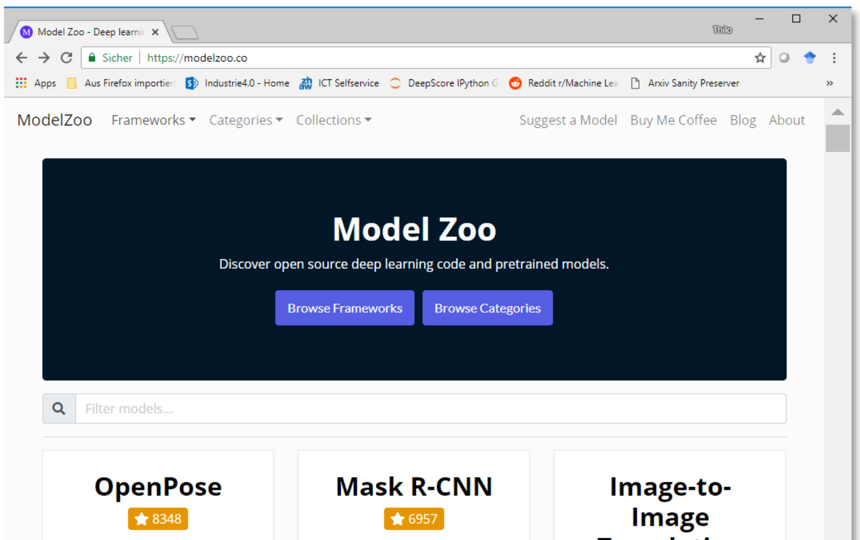

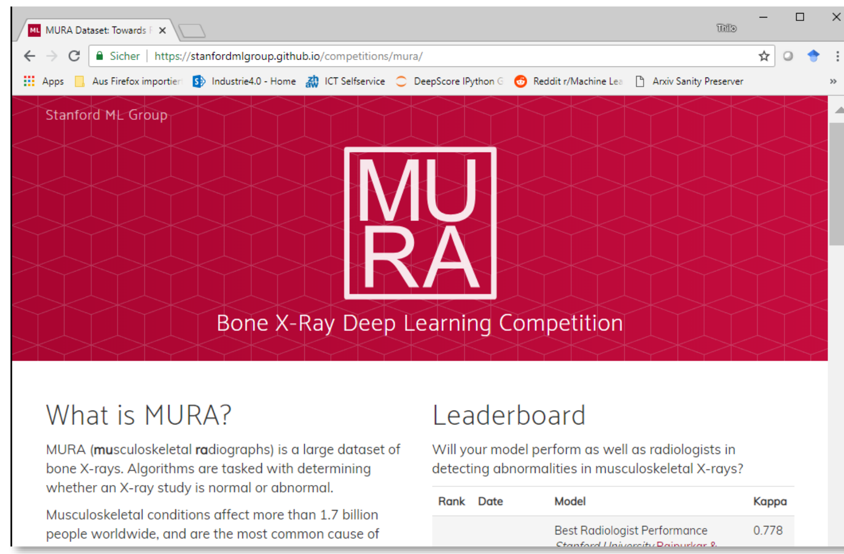

4. Music scanning

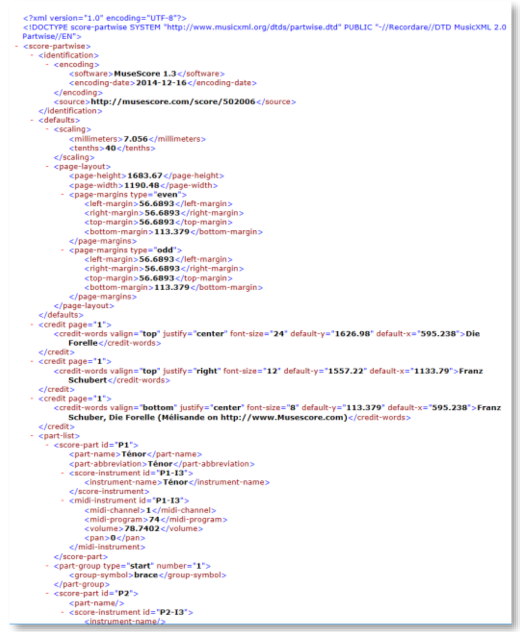

```
<?xml version="1.0" encoding="UTF-8"?>
<!DOCTYPE score-partwise SYSTEM "http://www.musicxml.org/dtds/partwise.dtd" PUBLIC "-//Recordare//DTD MusicXML 2.0 Partwise//EN">
- <score-partwise>
  - <identification>
    - <encoding>
        <software>MuseScore 1.3</software>
        <encoding-date>2014-12-16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


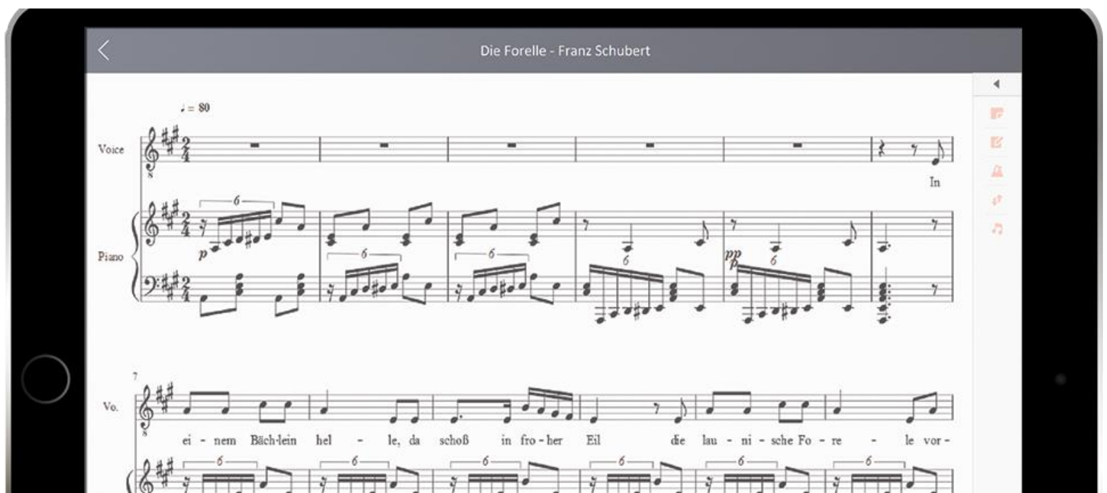

Zürcher Hochschule
für Angewandte Wissenschaften

Schweizerische Eidgenossenschaft
Confédération suisse
Confederazione Svizzera
Confederaziun svizra

Swiss Confederation

Innosuisse – Swiss Innovation Agency

4. Music scanning – challenges & solutions

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

Zürcher Hochschule
für Angewandte Wissenschaften

4. Music scanning – challenges & solutions

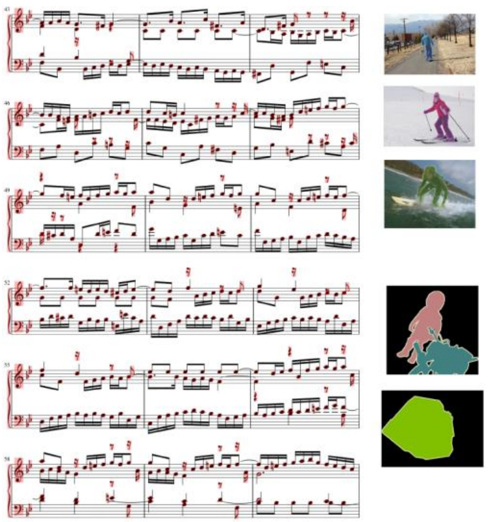

Schweizerische Eidgenossenschaft
Confédération suisse
Confederazione Svizzera
Confederaziun svizra

Swiss Confederation

Innosuisse – Swiss Innovation Agency

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

4. Music scanning – challenges & solutions

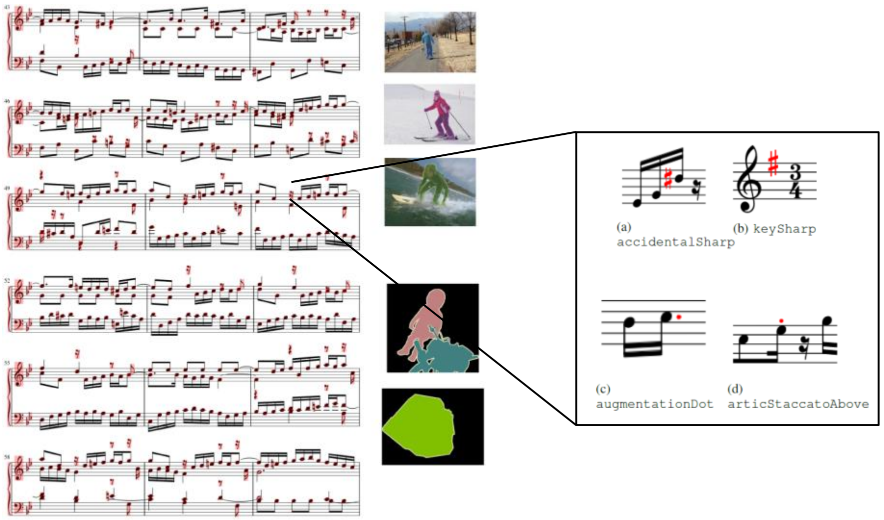

Schweizerische Eidgenossenschaft
Confédération suisse
Confederazione Svizzera
Confederaziun svizra
Swiss Confederation
Innosuisse – Swiss Innovation Agency

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

4. Music scanning – challenges & solutions

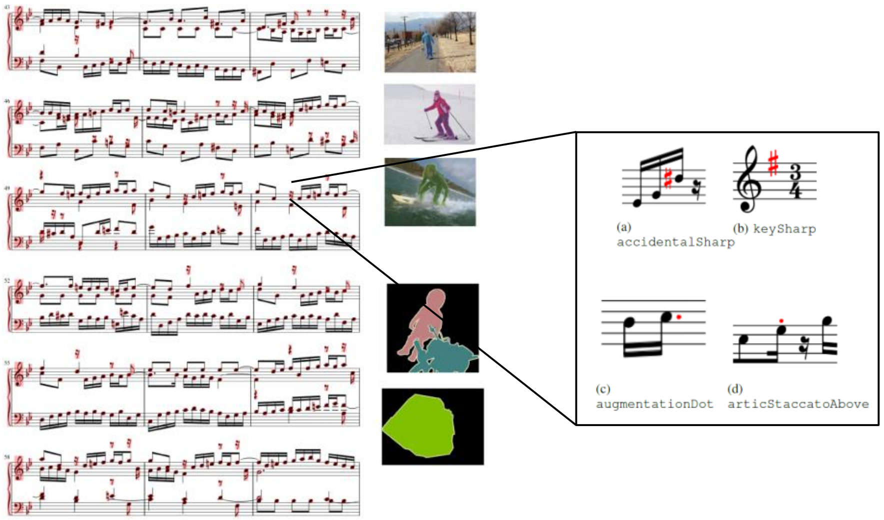

,

SCOREPAD

Schweizerische Eidgenossenschaft
Confédération suisse
Confederazione Svizzera
Confederaziun svizra

Swiss Confederation

Innosuisse – Swiss Innovation Agency

Tuggener, Elezi, Schmidhuber, Pelillo & Stadelmann (2018). *«DeepScores – A Dataset for Segmentation, Detection and Classification of Tiny Objects»*. ICPR’2018.

4. Music scanning – challenges & solutions

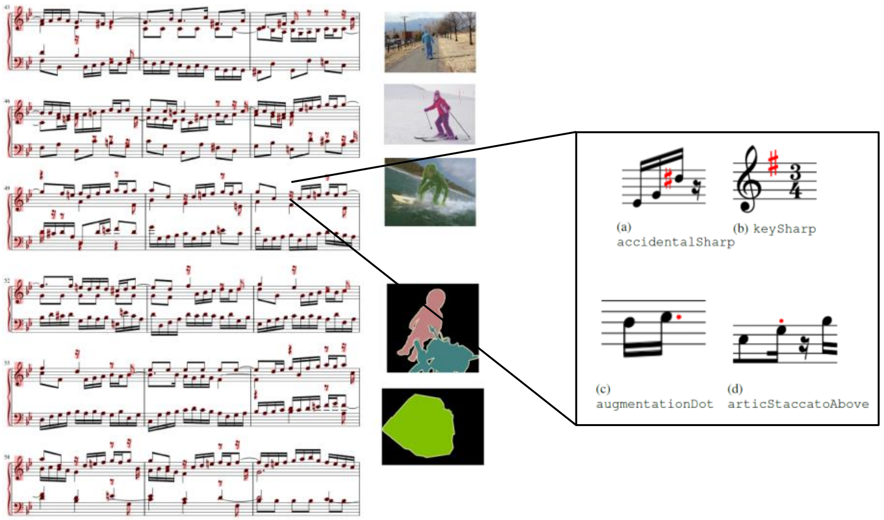

,

SCOREPAD

Schweizerische Eidgenossenschaft
Confédération suisse
Confederazione Svizzera
Confederaziun svizra

Swiss Confederation

Innosuisse – Swiss Innovation Agency

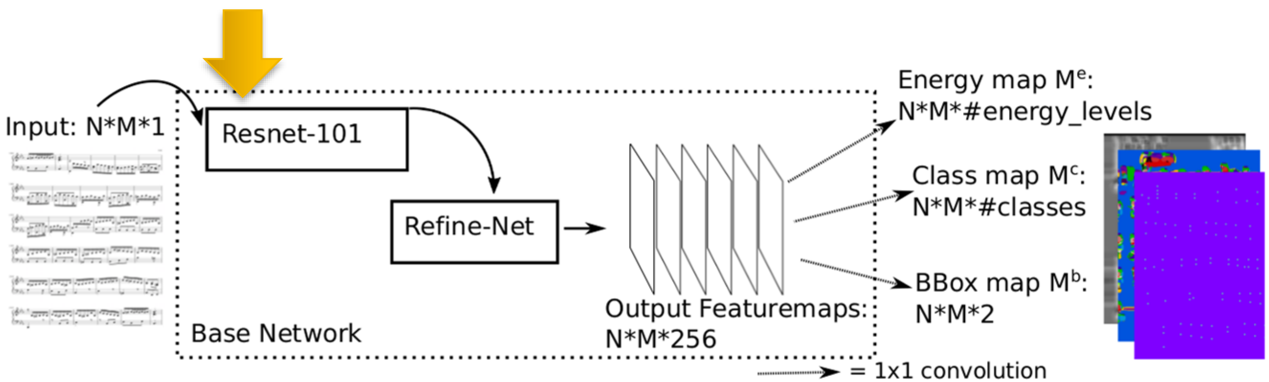

Tuggener, Elezi, Schmidhuber, Pelillo & Stadelmann (2018). *«DeepScores – A Dataset for Segmentation, Detection and Classification of Tiny Objects»*. ICPR'2018.
Tuggener, Elezi, Schmidhuber & Stadelmann (2018). *«Deep Watershed Detector for Music Object Recognition»*. ISMIR'2018.

4. Music scanning – challenges & solutions

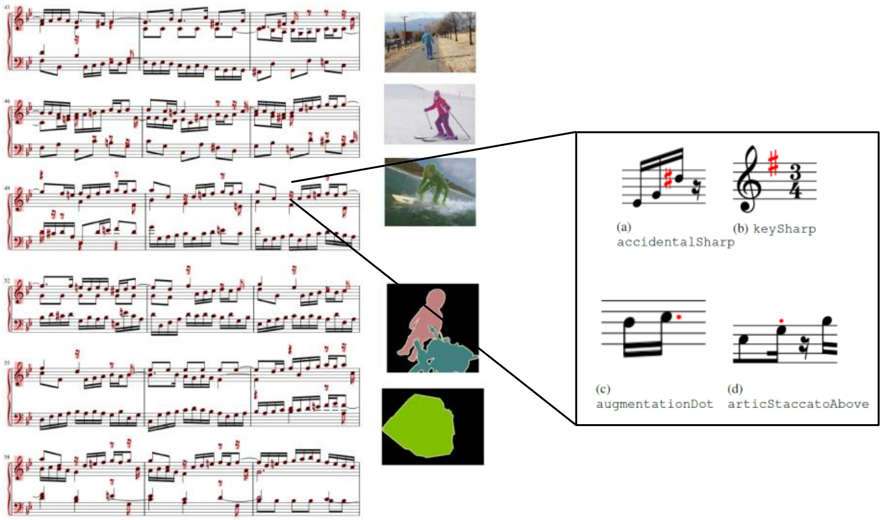

,

Schweizerische Eidgenossenschaft
Confédération suisse
Confederazione Svizzera
Confederaziun svizra

Swiss Confederation

Innosuisse – Swiss Innovation Agency

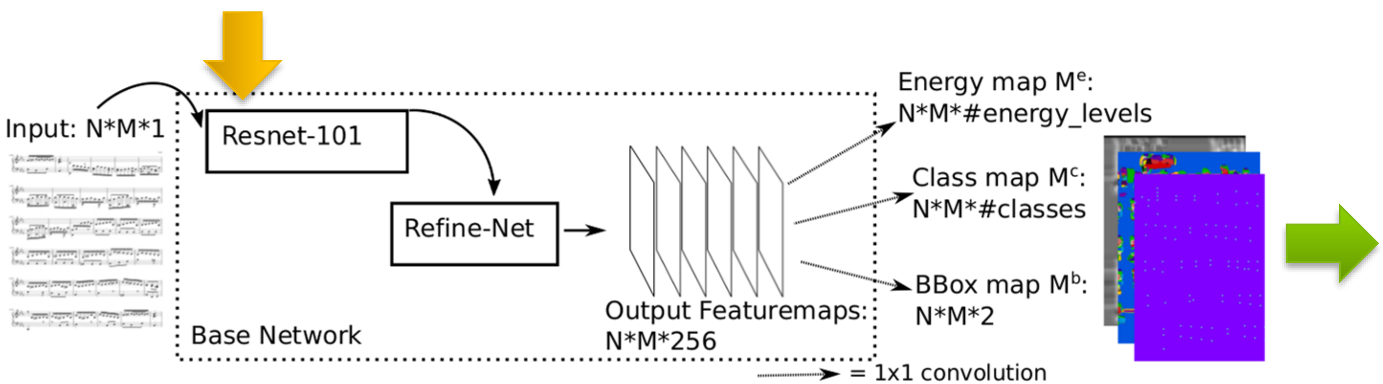

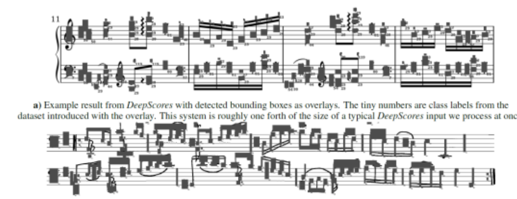

a) Example result from *DeepScores* with detected bounding boxes as overlays. The tiny numbers are class labels from the dataset introduced with the overlay. This system is roughly one forth of the size of a typical *DeepScores* input we process at once.

b) Example result from *MUSCIMA++* with detected bounding boxes and class labels as overlays. This system is roughly one half of the size of a typical processed *MUSCIMA++* input. The images are random picks amongst inputs with many symbols.

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.

4. Music scanning – industrialization
(Work in progress)

Recent results on class imbalance and robustness challenges

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

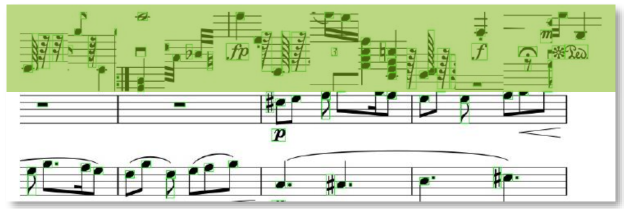

2. Put additional effort (and compute) into hyperparameter **tuning** and **longer training**
3. Trained also on scanned (more **real-worldish**) scores

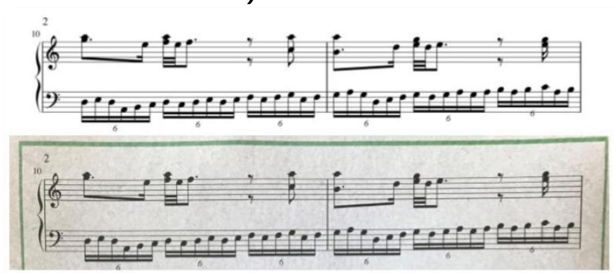

➔ **Improved** our **mAP** from 16% (on purely synthetic data) **to 73%** on more challenging real-world data set (additionally, using Pacha et al.'s evaluation method as a 2nd benchmark: from 24.8% to 47.5%)

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.

Lessons learned 4/4

, 

Loss shaping

- Usually necessary to **enable learning** of very complex target functions

“Initially, the training was ***unstable*** *[…] if directly trained on the* ***combined weighted loss****.* ***Therefore****, we now* ***train*** *[…] on each of the* ***three tasks separately****. We further observed that while the network gets trained on the bounding box prediction and classification, the energy level predictions get worse. To avoid this, the network is* ***fine-tuned only for the energy level loss*** *[…]. Finally, the network is retrained on the combined task […] for a few thousand iterations […].”*

- This includes **encoding expert knowledge** manually into the model architecture or training setup

“The ***size of the anomaly*** *in classifying balloon catheters as good or bad is* ***quite decisive****. Thus, rescaling the training images is not allowed, and we used a fixed size window around the center of each defect to extract the training images.”*

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

Hypothesis: basic & applied research converge

Speed of "digital" innovation makes complementary skills necessary *at the same time*:

- *Rigor* to come up with completely new methodical approaches
- *Creativity* to solve completely new scenario, thereby "filling wholes"

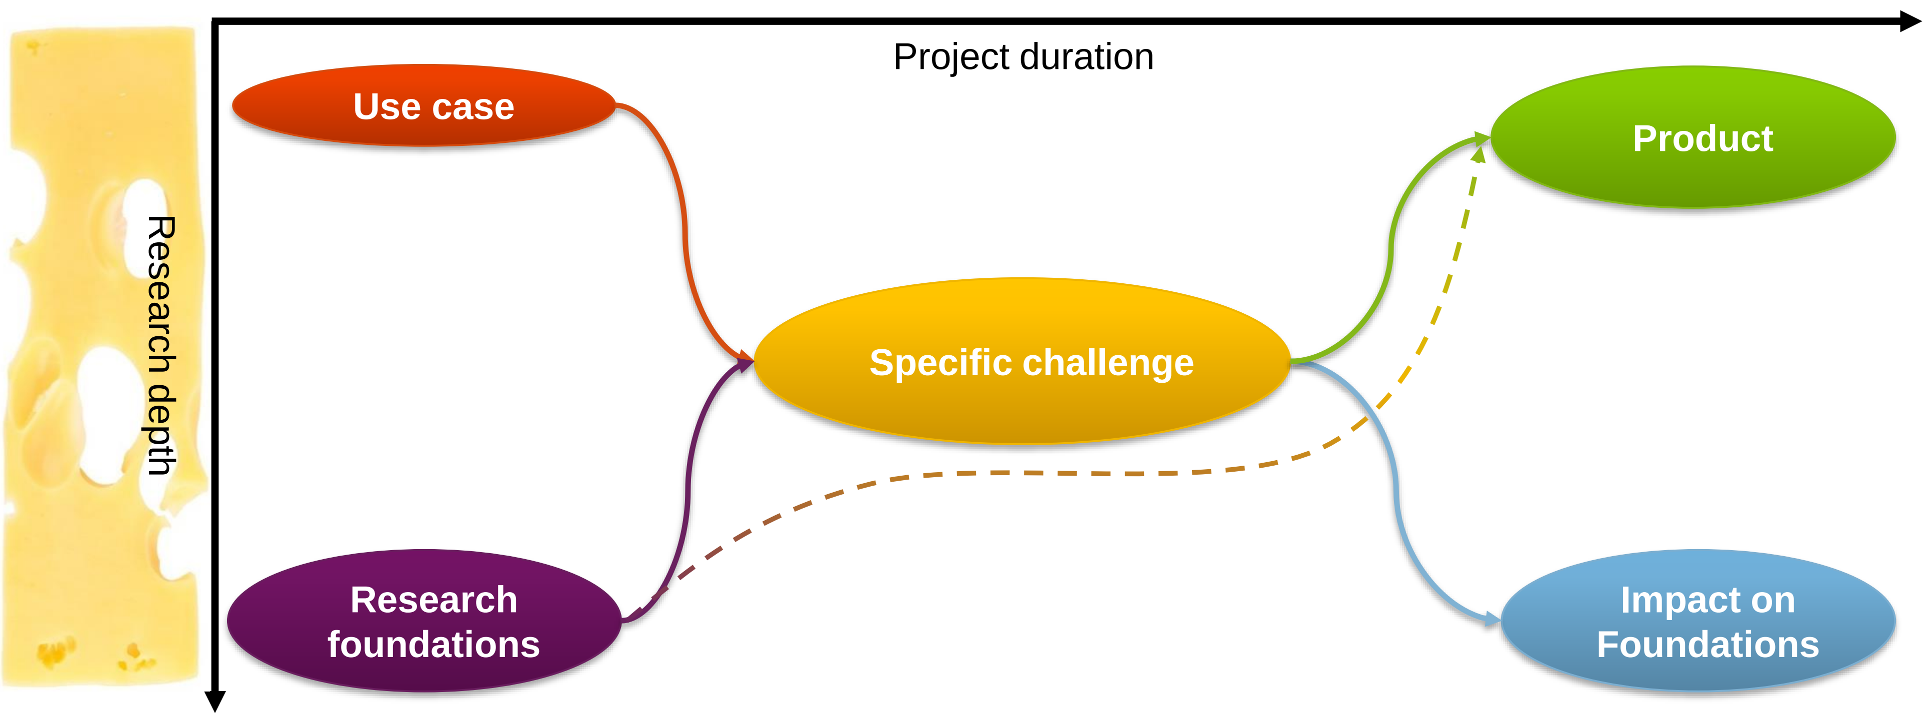

Conclusions

- **Latest research** is **applied** and deployed in «normal» businesses (non-AI, SME)
- It does not need big-, but some **data** (**effort usually underestimated**)
- DL/RL **training** for new use cases **can be tricky** (→ needs thorough experimentation)
- The **simultaneity** of research **types A$^{\text{pplied}}$ and B$^{\text{asic}}$** speaks out loud for **collaboration**

datalab
www.zhaw.ch/datalab

swiss group for artificial intelligence and cognitive science

SGAICO

data+service

Swiss Alliance for Data-Intensive Services

On me:

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

On the topics:

- AI: https://sgaico.swissinformatics.org/
- Data+Service Alliance: www.data-service-alliance.ch
- Collaboration: datalab@zhaw.ch

→ Happy to answer questions & requests.

APPENDIX

Zürcher Hochschule
für Angewandte Wissenschaften

5. Game playing
(work in progress)

GIANTS SOFTWARE

Schweizerische Eidgenossenschaft
Confédération suisse
Confederazione Svizzera
Confederaziun svizra

Swiss Confederation

Innosuisse – Swiss Innovation Agency

(symbolic figure)

5. Game playing – challenges & solutions
(work in progress)

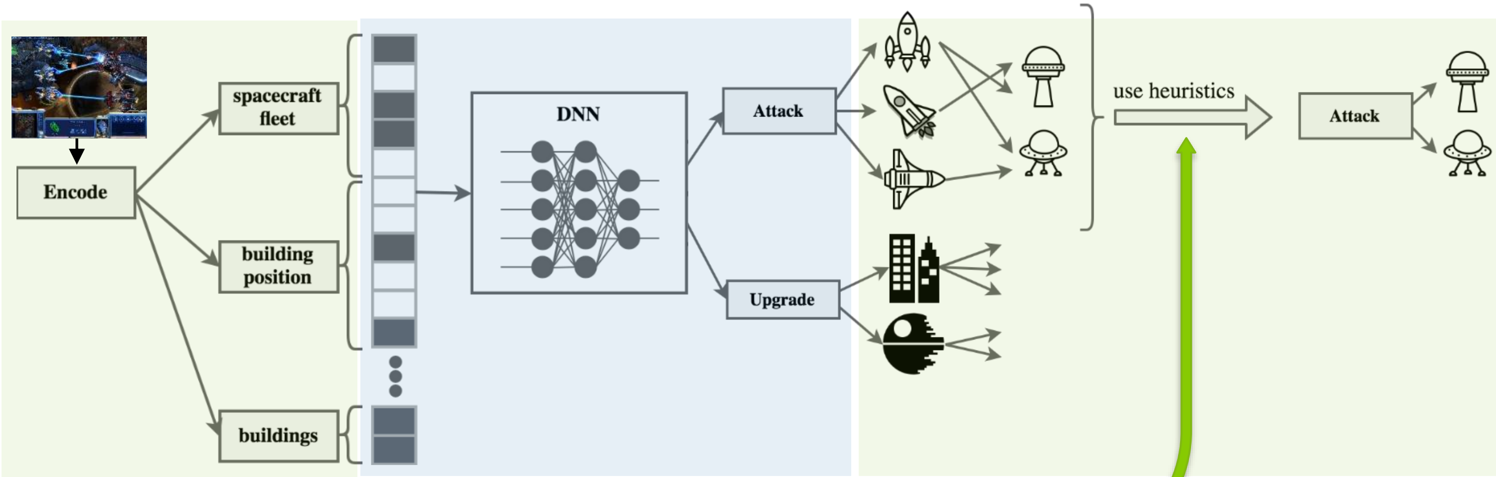

Reinforcement learning: deep Q network

Large discrete action space → use heuristic
- makes exploration difficult
- elongates training time

Delayed and sparse reward → do reward shaping
- sequence of actions crucial to get a reward

Distance encoding → use reference points

Transfer Learning → difficult: more complex environment needs other action sequence

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

6. Automated machine learning
(work in progress)

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

The project

- Target: in-house solution of industrial partner to improve turnover in standard analytics projects
- Challenge: optimize hyperparameters smarter than with well initialized random perturbations
- Idea: use reinforcement learning to meta-learn from past analytics projects

Initial experiments

			Auto-Sklearn		TPOT		DSM	
Dataset	**Task**	**Metric**	**Validation**	**Test**	**Validation**	**Test**	**Validation**	**Test**
Cadata	Regression	Coefficient Of Determination	0.7913	0.7801	**0.8245**	**0.8017**	0.7078	0.7119
Christine	Binary Classification	Balanced Accuracy Score	0.7380	0.7405	**0.7435**	**0.7454**	0.7362	0.7146
Digits	Multiclass Classification	Balanced Accuracy Score	**0.9560**	**0.9556**	0.9500	0.9458	0.8900	0.8751
Fabert	Multiclass Classification	Accuracy Score	**0.7245**	**0.7193**	0.7172	0.7006	0.7112	0.6942
Helena	Multiclass Classification	Balanced Accuracy Score	**0.3404**	**0.3434**	0.2654	0.2667	0.2085	0.2103
Jasmine	Binary Classification	Balanced Accuracy Score	0.7987	0.8348	**0.8188**	0.8281	0.8020	**0.8371**
Madeline	Binary Classification	Balanced Accuracy Score	**0.8917**	**0.8769**	0.8885	0.8620	0.7707	0.7686
Philippine	Binary Classification	Balanced Accuracy Score	0.7787	0.7486	**0.7839**	**0.7646**	0.7581	0.7406
Sylvine	Binary Classification	Balanced Accuracy Score	0.9414	0.9454	**0.9512**	**0.9493**	0.9414	0.9233
Volkert	Multiclass Classification	Accuracy Score	**0.7174**	**0.7101**	0.6429	0.6327	0.5220	0.5153
Average Performance			**0.7678**	**0.7654**	0.7586	0.7497	0.7048	0.6991

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

7. Condition monitoring
Maintaining machines on predicted failure only

We 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.

Mechmine data science technologies

Schweizerische Eidgenossenschaft
Confédération suisse
Confederazione Svizzera
Confederaziun svizra

Swiss Confederation

Innosuisse – Swiss Innovation Agency

vibration sensors

feature extraction

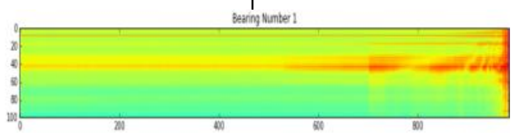

e.g., RNN autoencoder

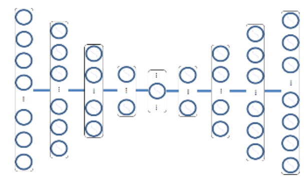

early detection of fault

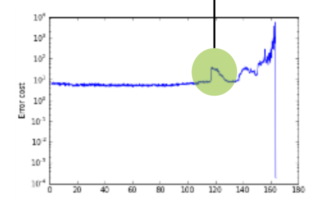

Stadelmann, Tolkachev, Sick, Stampfli & Dürr (2018). «*Beyond ImageNet - Deep Learning in Industrial Practice*». In: Braschler et al., «Applied Data Science», Springer.

8. Trace & detect adversarial attacks
…using average local spatial entropy of feature response maps

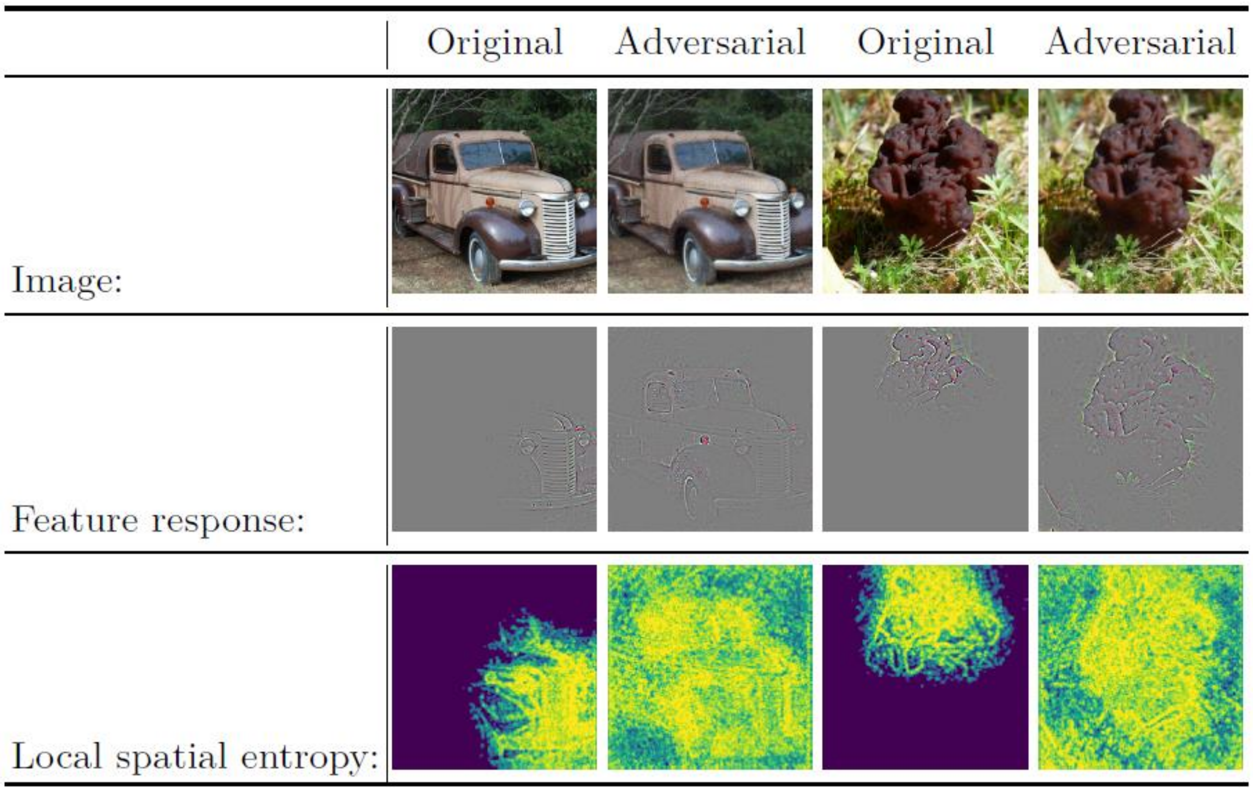

Amirian, Schwenker & Stadelmann (2018). *«Trace and Detect Adversarial Attacks on CNNs using Feature Response Maps»*. ANNPR'2018.

ML @ Information Engineering Group

Institute of Applied Information Technology, ZHAW School of Engineering

Zürcher Hochschule für Angewandte Wissenschaften

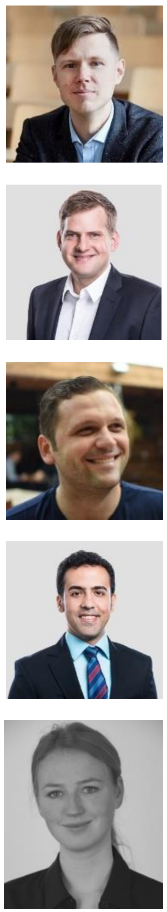

Machine learning-based Pattern Recognition

Robust Deep Learning

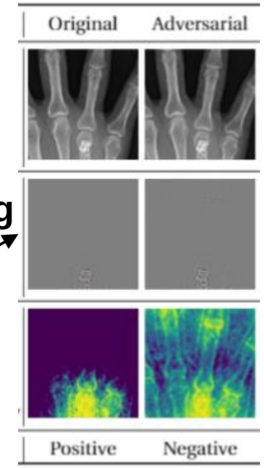

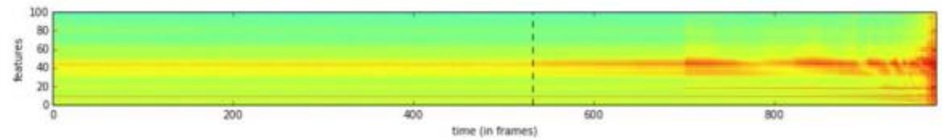

Voice Recognition

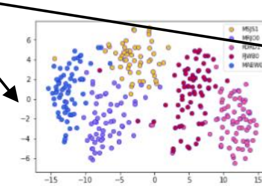

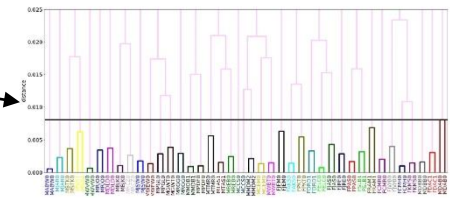

Document Analysis

Learning to Learn & Control