Lessons Learned Today for AI of Tomorrow

PwC Data Analytics All Hands, October 01, 2021

Thilo Stadelmann

Zürcher Hochschule
für Angewandte Wissenschaften

Zürcher Hochschule
für Angewandte Wissenschaften

Agenda

1. The new ZHAW
Centre for AI

2. Lessons Learned
from Current AI
Challenges

3. A Comprehensive
Vision for Developing
Tomorrow's AI

4. Discussion

The ZHAW Centre for Artificial Intelligence

Foundation: Machine Learning & Deep Learning
Cross-cutting: Ethics, society, more general AI

Autonomous Learning Systems

- *Reinforcement Learning*
- *Multi-Agent Systems*
- *Embodied AI*

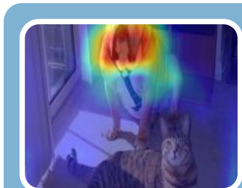

Explainable AI

- *Trustworthy Machine Learning*
- *Robust Deep Learning*
- *MLOps*

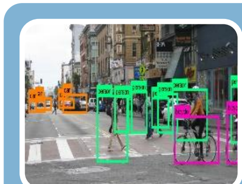

Computer Vision, Perception and Cognition

- *Pattern Recognition*
- *Machine Perception*
- *Neuromorphic Engineering*

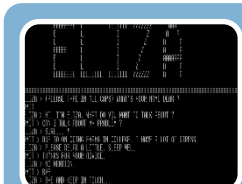

Natural Language Processing

- *Dialogue Systems*
- *Text Analytics*
- *Spoken Language Technologies*

Areas of application & cooperation: medicine & health, IoT, robotics, AI ethics & regulation, predictive maintenance, automatic quality control, document analysis, chat bots, biometrics, earth observation, digital farming, meteorology, autonomous driving, further data science use cases in industries like manufacturing / finance / insurance / commerce / transportation / energy etc.

Computer Vision, Perception & Cognition Group

Zürcher Hochschule für Angewandte Wissenschaften

Machine learning-based Pattern Recognition

Robust applications

[1]

[2]

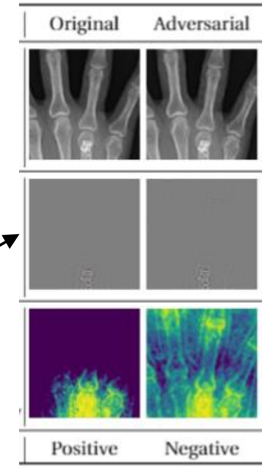

[3]

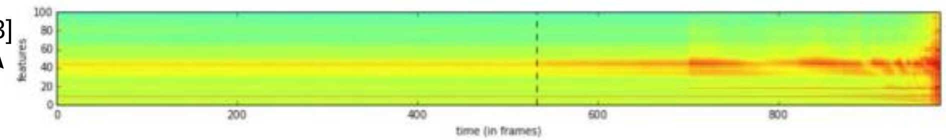

Biometrics

[4]

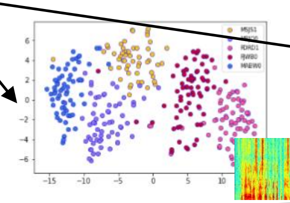

[5]

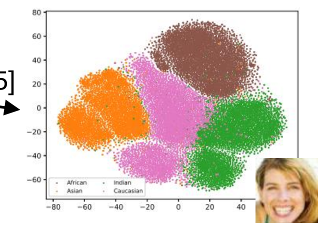

Document Analysis

[7]

[6]

Learning to act

[1]

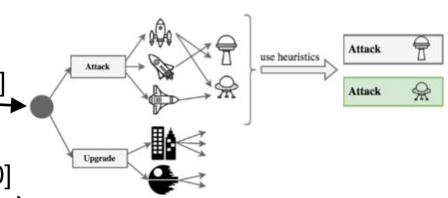

[8]

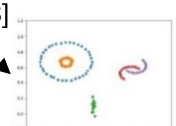

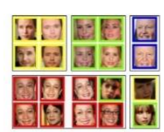

[9]

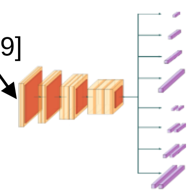

[10]

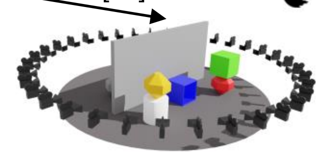

CVPC Group: recent results

- Medical image analysis: learning to reduce motion artifacts in 3D CT scans

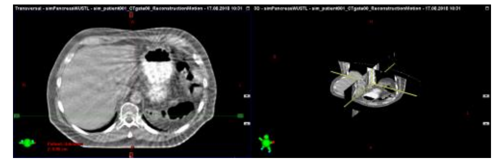

- Learning an artificial communication language for multi-agent reinforcement learning in logistics
  (notable rank in Flatland 2019 competition, best poster award [1])

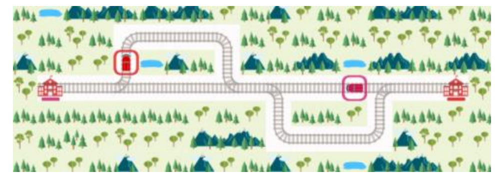

- Automated deep learning
  (benchmarked DSM [2], top rank in AutoDL 2020 challenge [3])

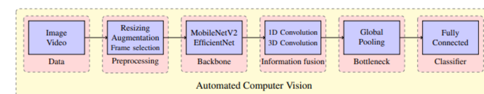

- Learning to segment and classify food waste in professional kitchens under adversarial conditions [5]

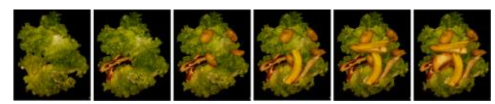

- Improving robotic vision through active vision and combined supervised and reinforcement learning
  (Dr. Waldemar Jucker Award 2020 [4])

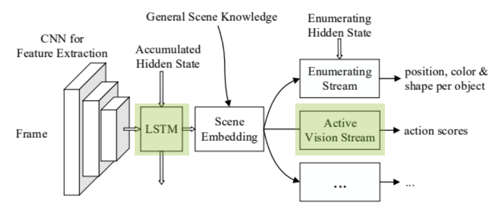

[1] Roost, Meier, Huschauer, Nygren, Egli, Weiler & Stadelmann (2020). *«Improving Sample Efficiency and Multi-Agent Communication in RL-based Train Rescheduling»*. SDS'2020.
[2] Tuggener, Amirian, Rombach, Lörwald, Varlet, Westermann & Stadelmann. *«Automated Machine Learning in Practice: State of the Art and Recent Results»*. SDS'2019.
[3] Tuggener, Amirian, Benites, von Däniken, Gupta, Schilling & Stadelmann (2020). *«Design Patterns for Resource Constrained Automated Deep Learning Methods»*. AI 1(4) 510-538.
[4] Roost, Meier, Toffetti Carughi & Stadelmann (2020). *«Combining Reinforcement Learning with Supervised Deep Learning for Neural Active Scene Understanding»*. AVHRC 2020.
[5] Simmler, Sager, Andermatt, Chavarriaga, Schilling, Rosenthal & Stadelmann (2021). *«A Survey of Un-, Weakly-, and Semi-Supervised Learning Methods for Noisy, Missing and Partial Labels in Industrial Vision Applications»*. SDS'2021.

CVPC Group: community outreach

Co-founder, **Swiss Conference Series on Data Science** (SDS)

- Largest annual Swiss event on data science, 400 participants in 2021
- Unique focus on business AND academia (IEEE technically co-sponsored)

Co-founder, **data innovation alliance**

- Innovation community of universities (19) and companies (56)
- Largest Swiss innovation network in the area of AI & data science

Initiator, 1st Int'l. **Symposium on the Science of Data Science**

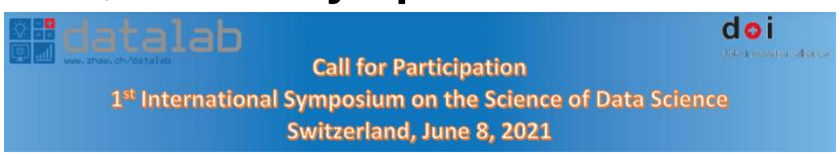

Book on lessons learned in applying AI & data science

- Braschler et al (eds.), "Applied Data Science – Lessons Learned for the Data-Driven Business", Springer, 2019

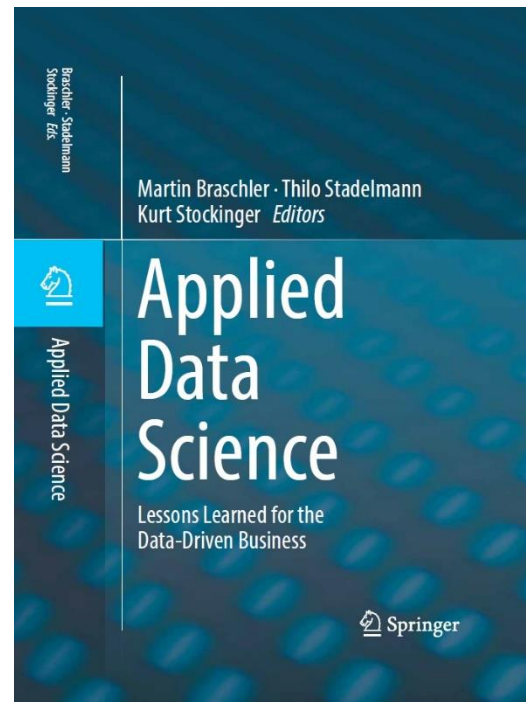

Zürcher Hochschule
für Angewandte Wissenschaften

Host of Swiss office of CLAIRE

Natural Language Processing Group

Chatbots
Chatbots can talk to customers and help them solve their tasks.

Dialogue Systems
We have the know-how to build reliable dialogue systems at scale.

Sentiment Analysis
Is a text positive or negative? Our algorithm won SemEval, one of the most prestigeous international competitions.

Topic Categorization
Distinguish hundreds of pre-defined topics/categories and label incoming documents accordingly.

Text-to-Speech
Generate audio from a given text in differente voices and speaker styles.

Automated Customer Support
Free up valuable resources without sacrificing the human bond to your customers.

Author Profiling
Segment your customer base with automatic age, gender, and geolocation detection from text messages.

Speech-to-Text
Automatically transcribe audio recordings into text.

Machine Translation
Tap into the power of machine learning to move seamlessly between different languages.

Summarization
Generate brief summaries of single documents or large document sets.

NLP Group: successfully completed projects

AuSUM
Automatic supply chainy monitoring
RepRisk - Innosuisse
1'200'000 CHF

Bio-SODA
Natural language interfaces to databases
Uni Lausanne, SIB - SNF
630'000 CHF

Call-E
Virtual agent for phone calls
Comparis - CTI
350'000 CHF

DIGI_KOM
Digital communication strategies
UAS Vorarlberg et al. - IBH
140'000 CHF

Speech Recognizer
Automatic speech recognition app
Sonova AG
50'000 CHF

SCAI
Smart contract analysis using AI
Legartis + CTI
910'000 CHF

Stiftungsregister SR
Foundation search engine
NonproCons - CTI
620'000 CHF

Community Cockpit
Cockpit for scalable customer interactions
Beekeeper, SpinningBytes - CTI
330'000 CHF

Speaking Robot
Speech-based demos for a robot head
School of Engineering (ZHAW)
138'000 CHF

Municipal Financing
Detection of financial needs of Swiss municipalities
SUVA
45'000 CHF

Interscriber
Transcription and summarization of dialogues
SpinningBytes AG - Innosuisse
800'000 CHF

DeepText
Intelligent text analysis with deep learning
SpinningBytes - CTI
500'000 CHF

KWS
Keyword spider
Eurospider - CTI
300'000 CHF

Headline Generation
Generation of teaser texts and headlines for news articles
Tamedia, SpinningBytes
62'000 CHF

Cleantech Cube
Detect cleantech products in company websites
Cleantech Switzerland
26'000 CHF

VirtualKids
Simulation of children in police interrogations
HSLU - SNF
790'000 CHF

PANOPTES
Newspaper auto-segmentation for live media monitoring
Argus Data Insights - CTI
490'000 CHF

DeLLA
Speech Recognizer with limited training data
SlowSoft AG - CTI
210'000 CHF

Talkalyzer
Share-in-Speech Analysis via Real-Time Speaker Classification
internal
50’000 CHF

Amazon AWS Grant
Grant for deep learning research
Amazon
22'000 CHF

LIHLITH
Lifelong learning for dialogue systems
UPV/EHU et al. - CHIST-ERA, SNF
680'000 CHF

Libra
One-tool solution for MLD4 compliance
DeepImpact - CTI
370'000 CHF

SwissText
Organization of SwissText conference, since 2016
CTI/Innosuisse
210.000 CHF

AUGEST
Automatic Generation of Regression Tests
Internal
50'000 CHF

Email Assistant
Concept for automatic email cleanup
Smart Data Way - CTI InnoScheck
7'500 CHF

Total Volume >8.5Mio CHF

NLP Group sample project: Interscriber

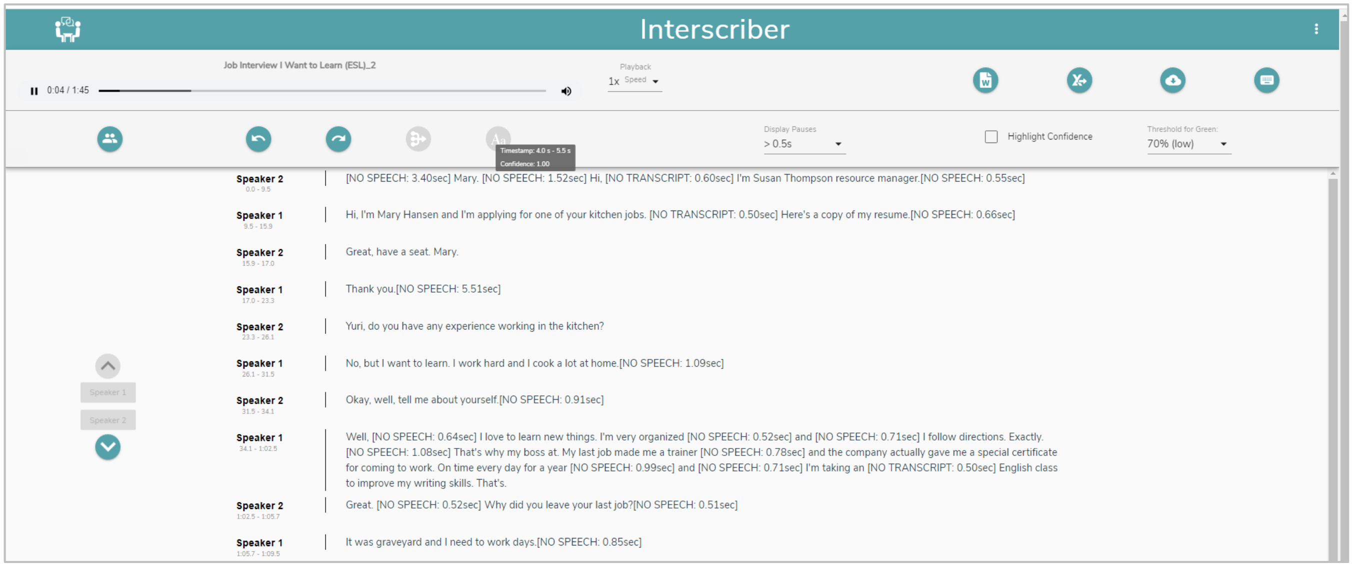

Interscriber (InnoSuisse, 420'000CHF)

- Automatic transcription of audio recordings to text (Speech-to-Text + Speaker Recognition)
- Smart editor with interactive player
- Automatic summaries; extraction of action points and generation of meeting minutes

NLP Group: active community contributor

- Organizer and Founder of the Swiss Text Analytics Conference (**SwisText**) in 2016, with more than 200 participants per year
- Co-Founder of the Swiss Association for Natural Language Processing (**SwissNLP**)
- 
  Publication of Corpora, e.g. for Sentiment Analysis and Speech Processing
- Organizer of Shared Tasks on Swiss German Recognition, Patent Classification etc.
- Co-Organizer of GermEval 2020
- Co-Organizer of Expert Group **NLP in Action**

2. Lessons Learned

Towards Robust Deep Learning, XAI, MLOps & Data-Centric AI

1. The new ZHAW
Centre for AI

2. Lessons Learned
from Current AI
Challenges

3. A Comprehensive
Vision for Developing
Tomorrow's AI

4. Discussion

Zürcher Hochschule
für Angewandte Wissenschaften

Robust Deep Learning

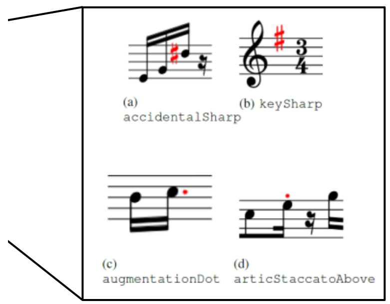

,

SCOREPAD

Schweizerische Eidgenossenschaft
Confédération suisse
Confederazione Svizzera
Confederaziun svizra

Swiss Confederation

Innosuisse – Swiss Innovation Agency

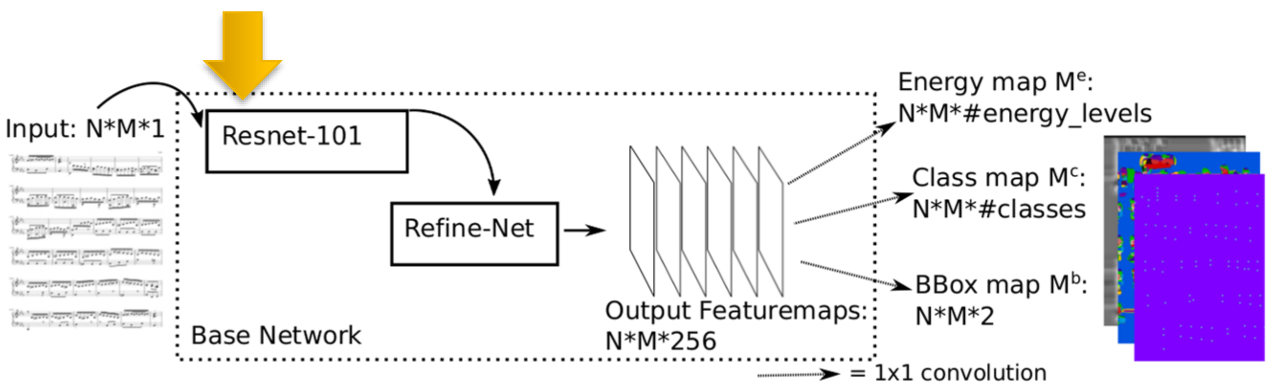

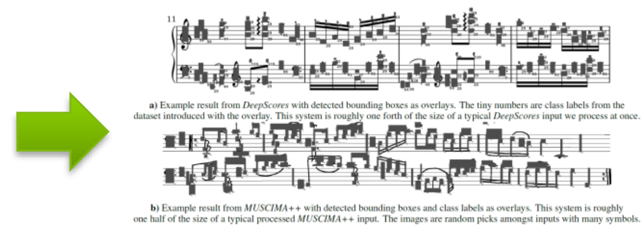

Tuggener, Elezi, Schmidhuber, Pelillo & Stadelmann (2018). *«DeepScores – A Dataset for Segmentation, Detection and Classification of Tiny Objects»*. ICPR'2018.
Tuggener, Elezi, Schmidhuber & Stadelmann (2018). *«Deep Watershed Detector for Music Object Recognition»*. ISMIR'2018.
Tuggener, Satyawan, Pacha, Schmidhuber & Stadelmann (2021). *«The DeepScoresV2 Dataset and Benchmark for Music Object Detection»*. ICPR'2020.

Robust Deep Learning
Industrialization

Current results on class imbalance and robustness challenges

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

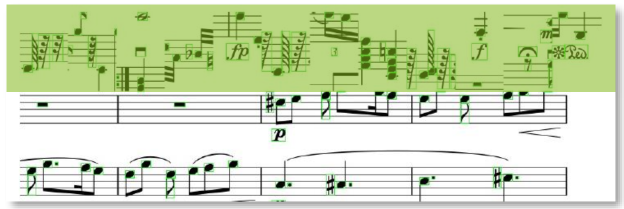

2. Put additional effort (and compute) into hyperparameter **tuning** and **longer training**
3. Trained also on scanned (more **real-worldish**) scores

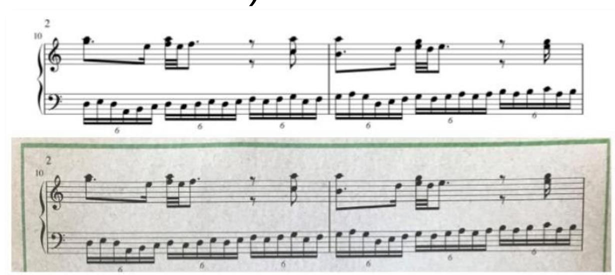

➔ **Improved** our **mAP** from 16% (on purely synthetic data) **to 73%** on more challenging real-world data set (additionally, using Pacha et al.'s evaluation method as a 2^{nd} benchmark: SotA from 24.8% to 47.5%)

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.

Robust Deep Learning

Dealing with Real-World Noise in Music Scanning

Zürcher Hochschule
für Angewandte Wissenschaften

Synthetic quality + labels

Real world quality

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

Printer and scanner artifacts

Wrinkles

Dirt

...

Model training → **Data distribution shift** → Model deployment

Remedy:
Use GANs and Domain Adaptation to disentangle representation from distribution

Explainable AI

Model interpretability matters in applications involving humans

Interpretability is required.

- 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

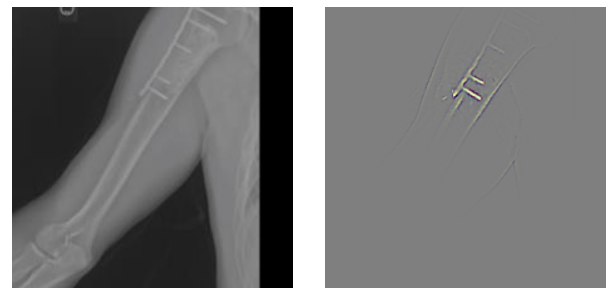

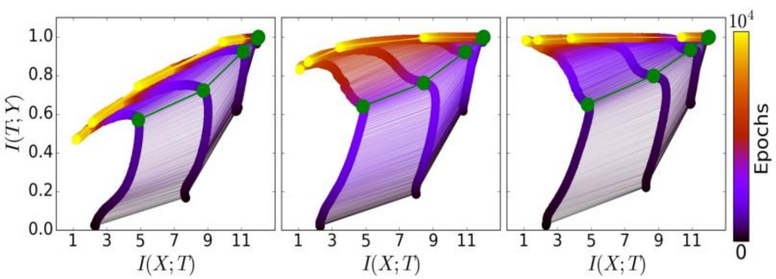

DNN training on the Information Plane

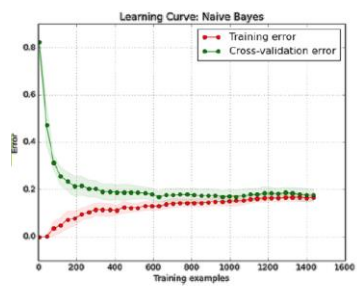

a learning curve

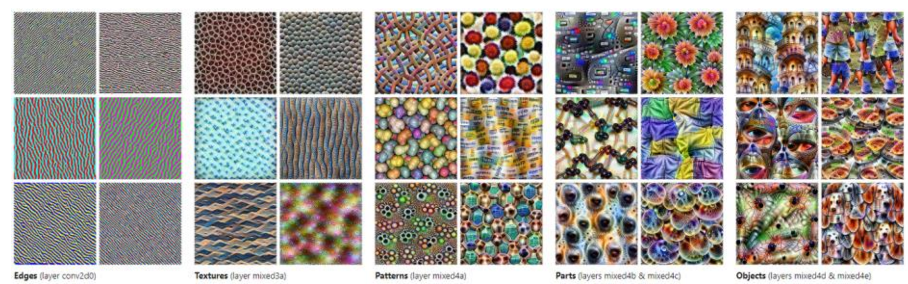

feature visualization

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/

Explainable AI

Detecting Adversarial Attacks…
…using average local spatial entropy of feature response maps

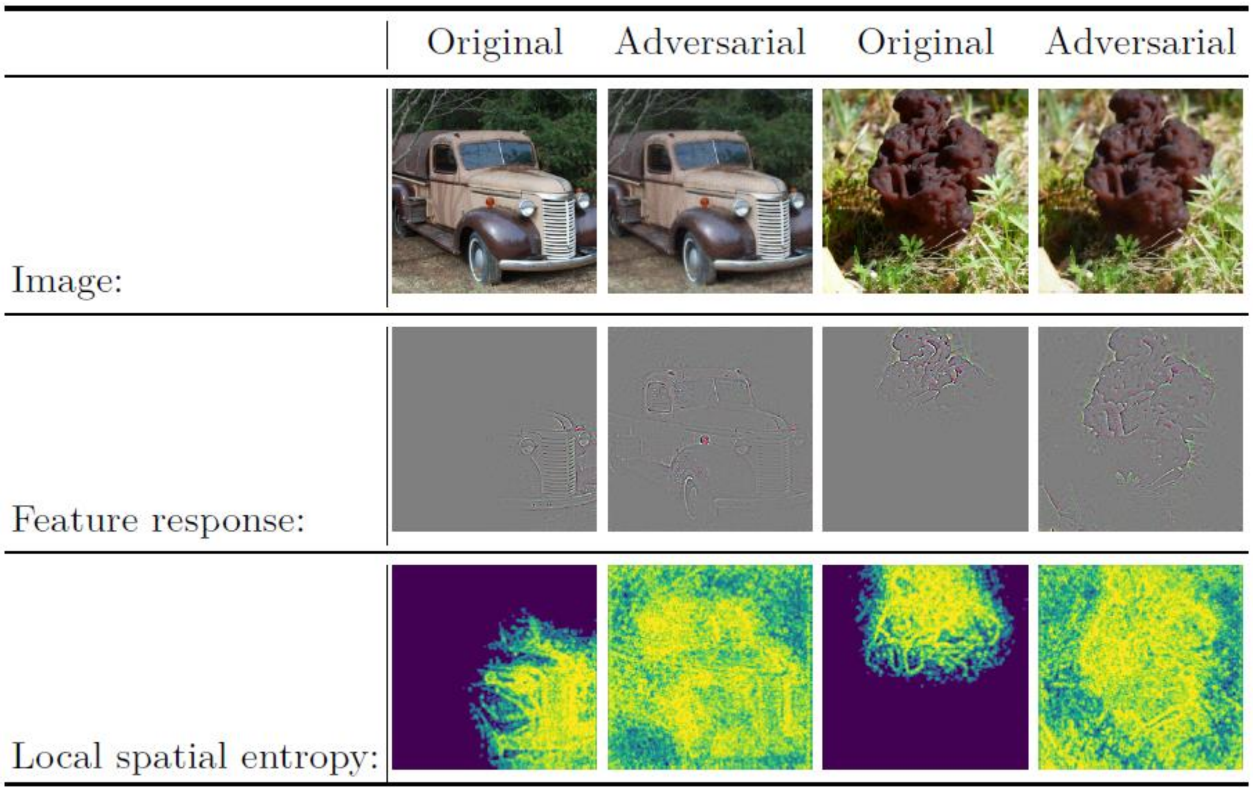

Amirian, Schwenker & Stadelmann (2018). *«Trace and Detect Adversarial Attacks on CNNs using Feature Response Maps»*. ANNPR'2018.

Zürcher Hochschule
für Angewandte Wissenschaften

MLOps

Two cases: Print Media Monitoring and Face Recognition

Task

Challenge

Nuisance

ARGUS DATA INSIGHTS
WISSEN ZUM ERFOLG

Schweizerische Eidgenossenschaft
Confédération suisse
Confederazione Svizzera
Confederaziun svizra

Swiss Confederation

Innosuisse – Swiss Innovation Agency

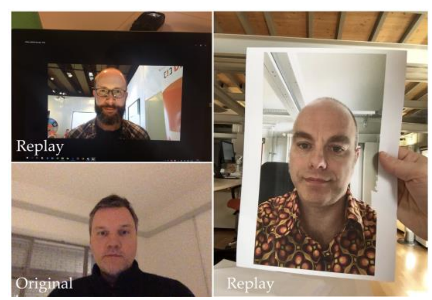

[!] DEEPIMPACT

Schweizerische Eidgenossenschaft
Confédération suisse
Confederazione Svizzera
Confederaziun svizra

Swiss Confederation

Innosuisse – Swiss Innovation Agency

MLOps
Complex pipelines need to be deployed, operated, maintained, …

ARGUS DATA INSIGHTS
WISSEN ZUM ERFOLG

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

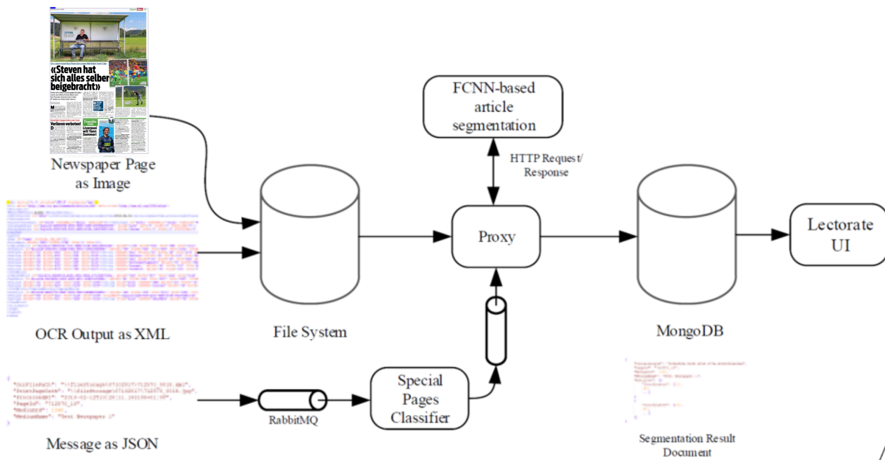

[!] DEEPIMPACT

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

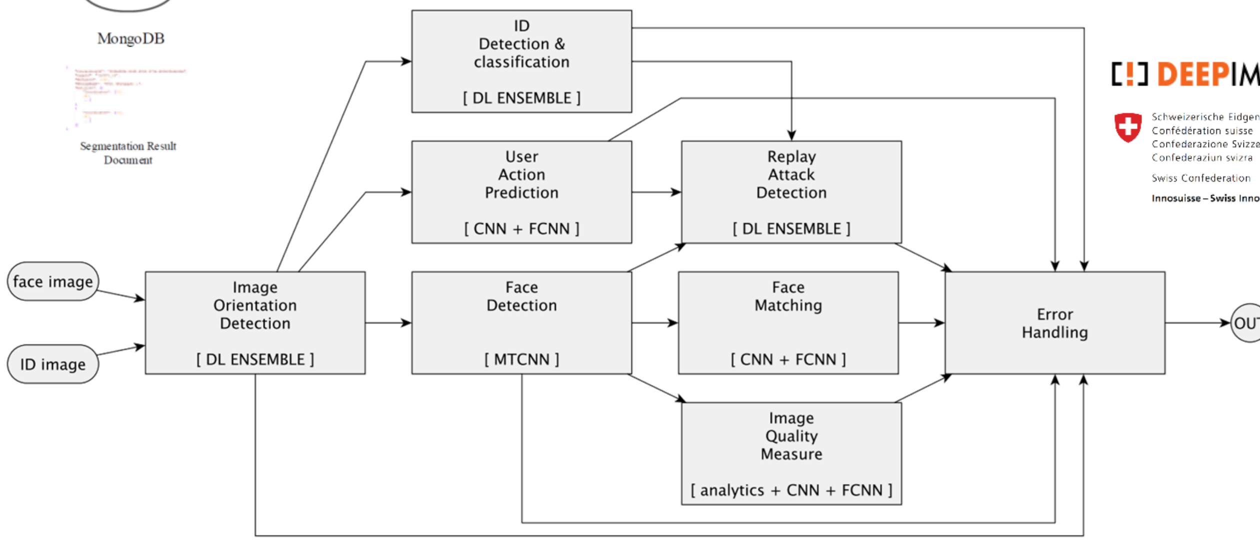

Stadelmann, Amirian, Arabaci, Arnold, Duivesteijn, Elezi, Geiger, Lörwald, Meier, Rombach & Tuggener (2018). *«Deep Learning in the Wild»*. ANNPR’2018.

Data-Centric AI
Two cases: Industrial Quality Control and Food Waste Analysis

Vastly varying defect sizes, poor label quality

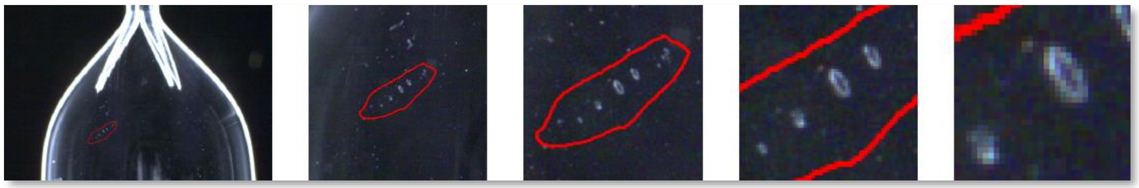

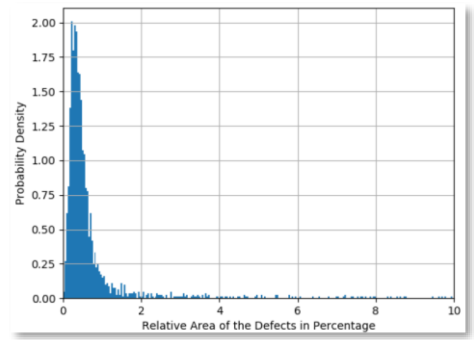

Vastly varying conditions, few labeled data, partial labels

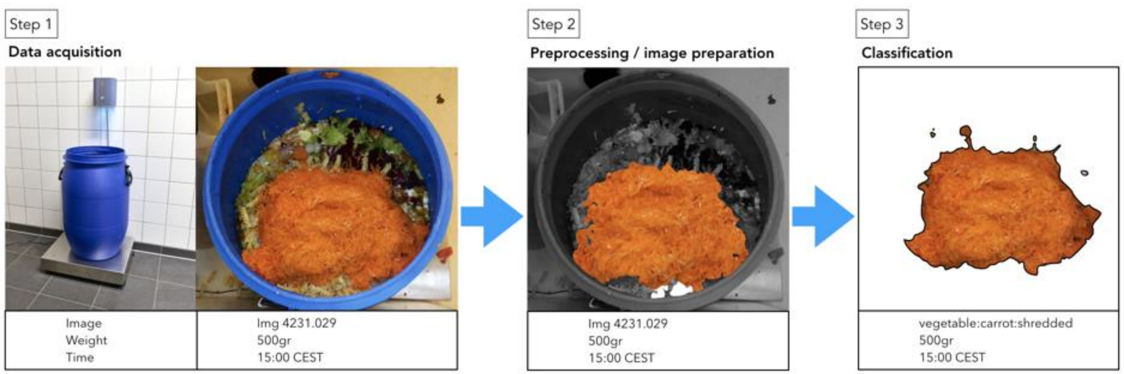

Data-Centric AI
Some lessons learned

Data needs more attention (as compared to modeling)
- Acquisition usually **needs much more time** than expected, yet is the basis for all subsequent success
- Class **imbalance** & **covariate shift** are typical, not abnormal

Unsupervised approaches need more attention

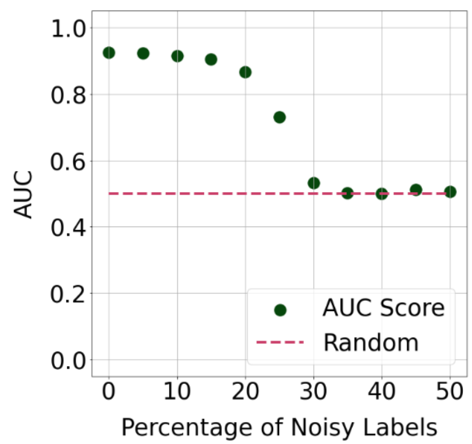

Simmler, Sager, Andermatt, Chavarriaga, Schilling, Rosenthal & Stadelmann (2021). *«A Survey of Un-, Weakly-, and Semi-Supervised Learning Methods for Noisy, Missing and Partial Labels in Industrial Vision Applications»*. SDS'2021.

3. Vision for Tomorrow’s AI

Zürcher Hochschule
für Angewandte Wissenschaften

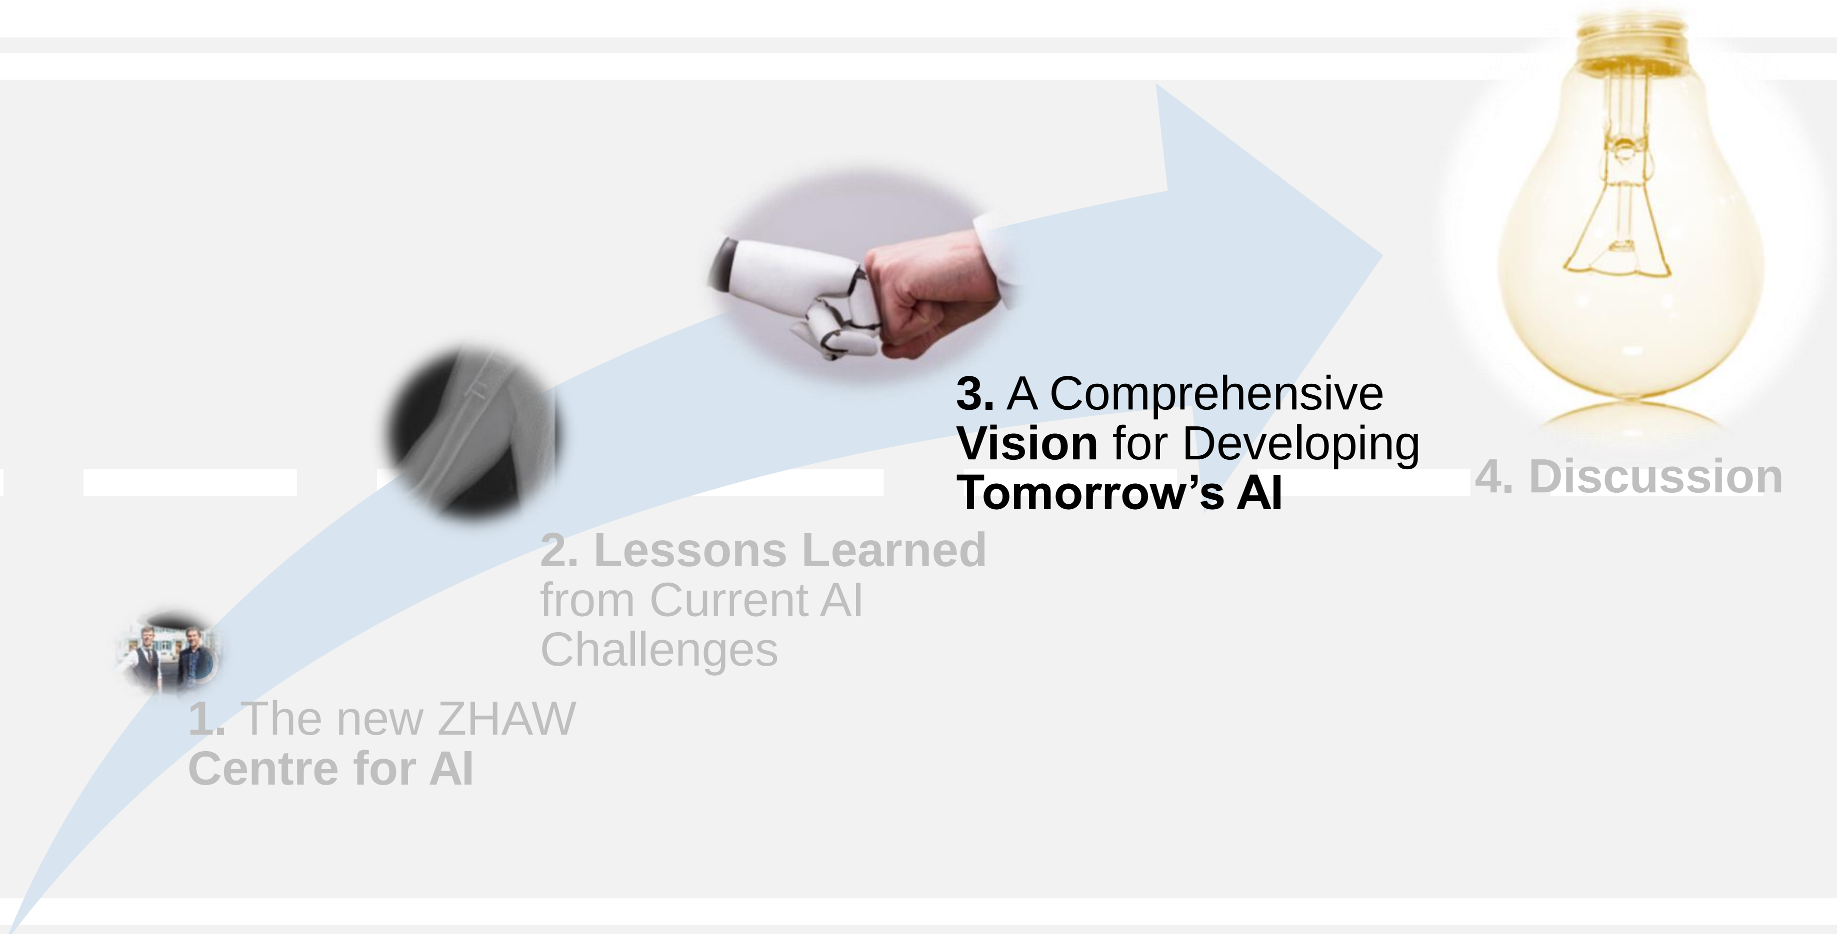

Human-AI Teaming provides learning opportunities to both partners (co-learning)

Successive stages of human-machine collaboration
1 (Mere tool): Human uses machine as mere UI
2 (Conventional computer vision): Trained CV system makes predictions that the human may consider (no learning interactions after training)
3 (Proposed co-learning): CV system continually learns from human preferences expressed in interactions; human learns from AI insights on own process

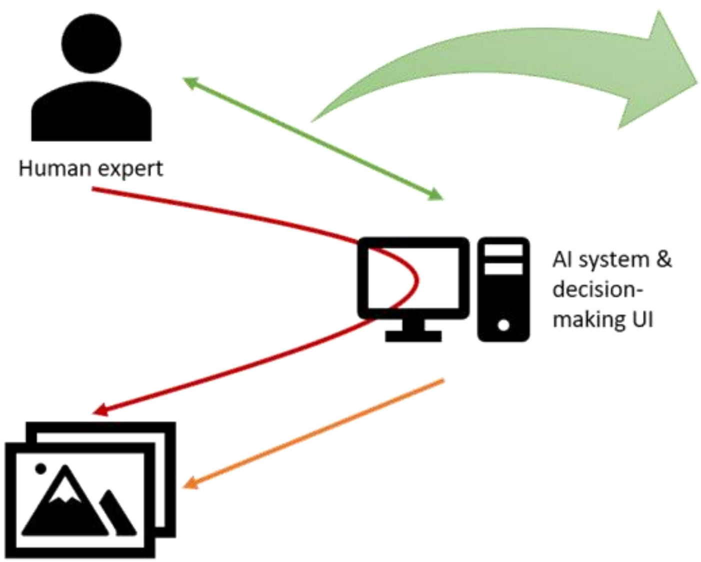

Non-obvious visual decision task, e.g.: medical diagnosis, surveillance / earth observation, complex expert reports

How to co-learn? High-level sources for mutual learning

1. AI system can adapt continually to human preferences using
   a) ...explicit corrections made to AI decisions (feedback) >> continual learning (active / transfer learning)
   b) ...implicit observations of human decision-making process >> contrastive learning, reinforcement learning
   c) ...explicit hints to analogous situations not considered similar by AI so far >> deep case-based reasoning
2. Human can learn from AI as it provides
   a) ...different features & local vision approach (complementary strengths) >> XAI
   b) ...hints from analyzing the human interaction with the system via UI >> behavior analysis, recommender systems
   c) ...quick tests of hypothesis by quickly pulling up examples from all cases seen so far >> graph-NN, image retrieval

Why to co-learning? Overarching goals for trustworthy AI

- Increased trust in AI system (through insights into inner workings [2a] and feedback mechanism [1a])
- Improved robustness of AI system (through continuous learning [1a-c])
- Maximized AI efficiency (through exploiting the complementary strengths of human & machine [1-2])
- Responsible decisions (through mutual de-biasing in joint decisions with mutual insight into the decision process [1a,2b])

How to co-learning, specifically? Scientific approach & challenges for AI system development

Deal with the two tasks (computer vision to arrive at a decision, give reasonable feedback to human based on interactions) in a unified AI architecture based on deep learning methodology:

- Core: multi-task (to account for 2 tasks) interpretable computer vision architecture (>> interpretable CNN, Grad-CAM etc.)
- World model: graph of relationships among cases seen so far defined by human preferences (>> graph-NN, deep CBR)
- Continual training loop: incorporate feedback while countering catastrophic forgetting (>> RL, contrastive learning)

Towards higher levels of intelligent behaviour in AI systems

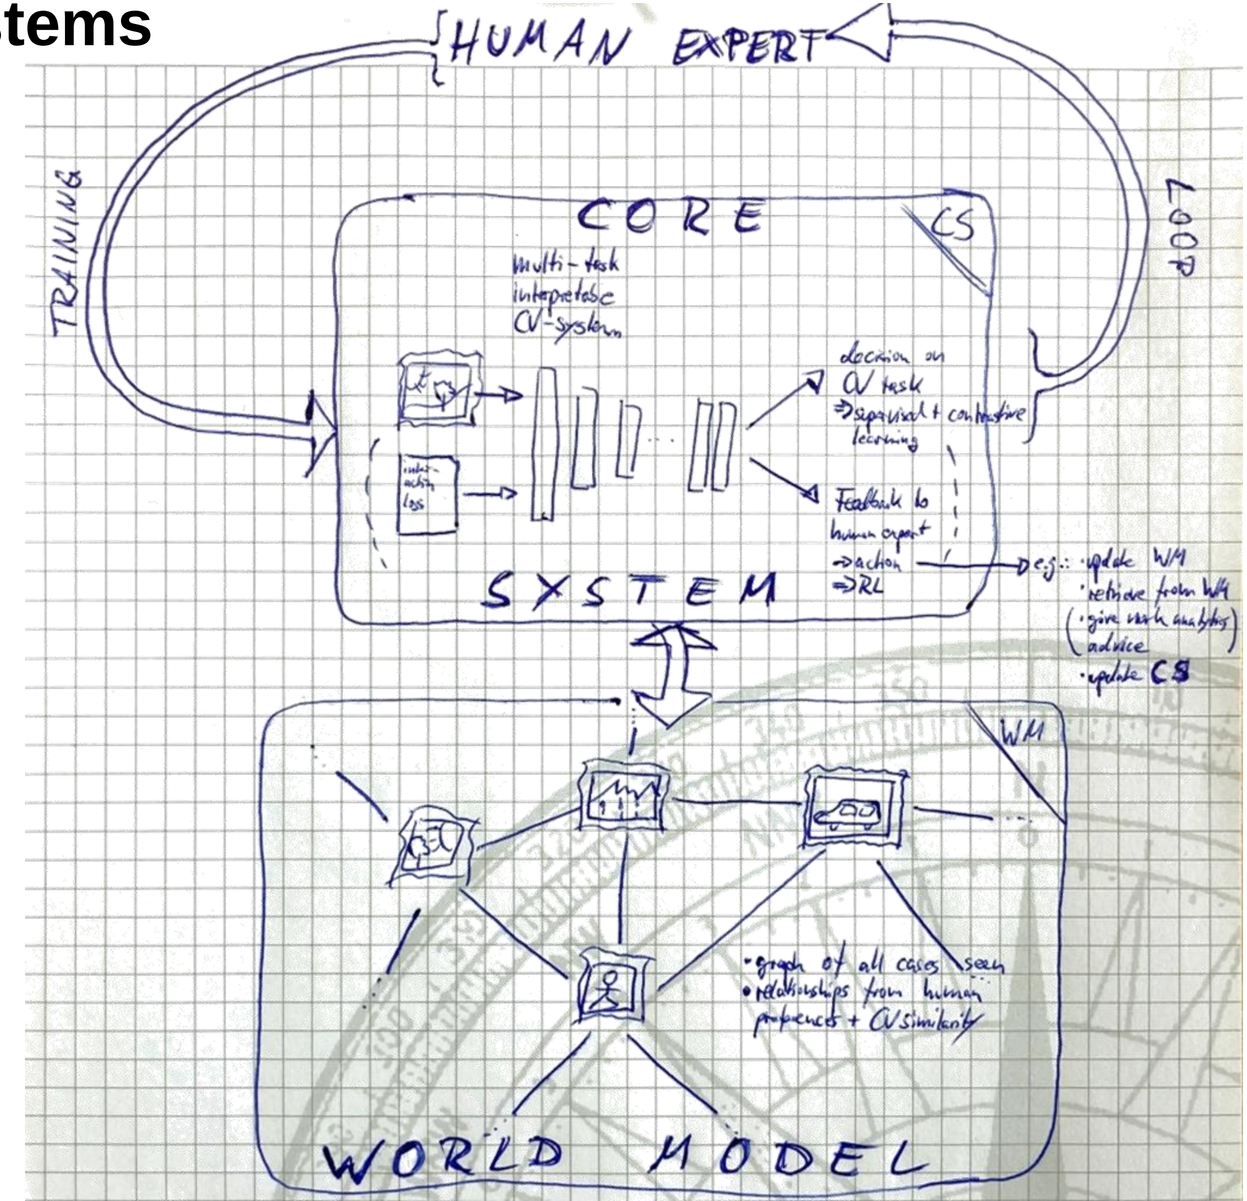

4. Discussion

- Interested in diverse use cases for machine learning and AI
- Work application-focused and methods-oriented
- Happy to collaborate interdisciplinary & internationally

About us:

- Director of Centre for AI, head CVPC Group: Prof. Dr. Thilo Stadelmann
  Email: stdm@zhaw.ch
  Phone: +41 58 934 72 08
- Head NLP Group: Prof. Dr. Mark Cieliebak
  Email: ciel@zhaw.ch
  Phone: +41 58 934 72 39

Further contacts:

- info.cai@zhaw.ch, datalab@zhaw.ch, info.office@data-innovation.org, office-switzerland@claire-ai.org

zhaw.ch/datalab

data innovation alliance

Swiss Association for
Natural Language Processing

APPENDIX

Sample projects

Why AI?


Can your phone improve your health?
Vitamin C and cancer cell metabolism
In search of stem cells
Science
11 DECEMBER 2015
AAAS
Machine intelligence
AI learns to read and write letters
nature
At last — a computer program that can beat a champion Go player
ALL SYSTEMS GO
SONGBIRDS À LA CARTE
SAFEGUARD TRANSPARENCY
arXiv monthly submission rates
12,000
11,000
10,000
9,000
8,000
7,000
6,000
5,000
4,000
3,000
2,000
1,000
0
Jan '92
Jan '14
Jan '16
Jan '18
IMAGENET
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
“The growth of deep-learning models is expected to accelerate and create even more innovative applications in the next few years.”

CAI Highlights

Education at the CAI

TEACHING ENGAGEMENT

- B.Sc. Computer Science & Data Science
- M.Sc. Engineering (CS, DS)
- Ph.D. in cooperation with e.g.

Universität Zürich[UZH] Università Ca'Foscari Venezia USI/SUPSI IDSIA Istituto Dalle Molle di studi sull'intelligenza artificiale universität ulm

- Continuing education in AI & ML
- Special mentoring program for CAI-affiliated students

UNDERGRAD PORTFOLIO

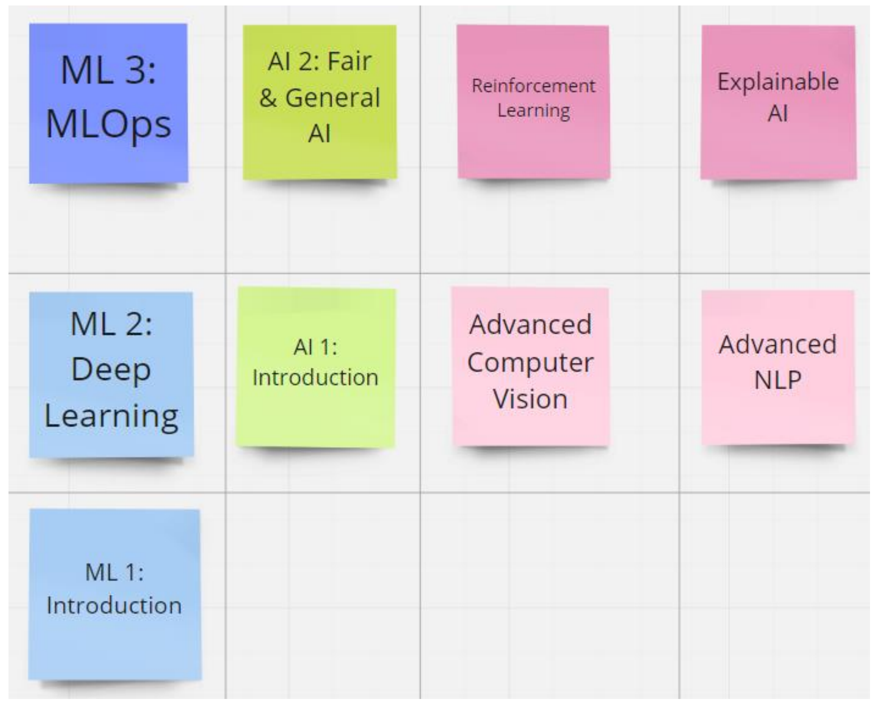

CVPC Group: references for overview

1. Thilo Stadelmann, Mohammadreza Amirian, Ismail Arabaci, Marek Arnold, Gilbert François Duivesteijn, Ismail Elezi, Melanie Geiger, Stefan Lörwald, Benjamin Bruno Meier, Katharina Rombach, and Lukas Tuggener. **"Deep Learning in the Wild"**. In: Proceedings of the 8th IAPR TC 3 Workshop on Artificial Neural Networks for Pattern Recognition (**ANNPR'18**), Springer, LNAI 11081, pp. 17-38, Siena, Italy, September 19-21, 2018.
2. Mohammadreza Amirian, Friedhelm Schwenker, and Thilo Stadelmann. **"Trace and Detect Adversarial Attacks on CNNs using Feature Response Maps"**. In: Proceedings of the 8th IAPR TC 3 Workshop on Artificial Neural Networks for Pattern Recognition (**ANNPR'18**), Springer, LNAI 11081, pp. 346-358, Siena, Italy, September 19-21, 2018.
3. Thilo Stadelmann, Vasily Tolkachev, Beate Sick, Jan Stampfli, and Oliver Dürr. **"Beyond ImageNet - Deep Learning in Industrial Practice"**. In: Martin Braschler, Thilo Stadelmann, and Kurt Stockinger (Editors). **"Applied Data Science - Lessons Learned for the Data-Driven Business"**. **Springer**, 2019.
4. Thilo Stadelmann, Sebastian Glinski-Haefeli, Patrick Gerber, and Oliver Dürr. **"Capturing Suprasegmental Features of a Voice with RNNs for Improved Speaker Clustering"**. In: Proceedings of the 8th IAPR TC 3 Workshop on Artificial Neural Networks for Pattern Recognition (**ANNPR'18**), Springer, LNAI 11081, pp. 333-345, Siena, Italy, September 19-21, 2018.
5. Stefan Glüge, Mohammadreza Amirian, Dandolo Flumini, and Thilo Stadelmann. **"How (Not) to Measure Bias in Face Recognition Networks"**. In: Proceedings of the 9th IAPR TC 3 Workshop on Artificial Neural Networks for Pattern Recognition (**ANNPR'20**), Springer, LNAI, Winterthur, Switzerland, September 02-04, 2020.
6. Lukas Tuggener, Yvan Putra Satyawan, Alexander Pacha, Jürgen Schmidhuber, and Thilo Stadelmann. **"The DeepScoresV2 Dataset and Benchmark for Music Object Detection"**. In: Proceedings of the 25th International Conference on Pattern Recognition (**ICPR'20**), IAPR, Milan, Italy, January 10-15 (online), 2021.
7. Benjamin Meier, Thilo Stadelmann, Jan Stampfli, Marek Arnold, and Mark Cieliebak. **"Fully convolutional neural networks for newspaper article segmentation"**. In: Proceedings of the 14th IAPR International Conference on Document Analysis and Recognition (**ICDAR'17**). 14th IAPR International Conference on Document Analysis and Recognition (ICDAR), Kyoto Japan, November 13-15, 2017. Kyoto, Japan: CPS.
8. Benjamin Bruno Meier, Ismail Elezi, Mohammadreza Amirian, Oliver Dürr, and Thilo Stadelmann. **"Learning Neural Models for End-to-End Clustering"**. In: Proceedings of the 8th IAPR TC 3 Workshop on Artificial Neural Networks for Pattern Recognition (**ANNPR'18**), Springer, LNAI 11081, pp. 126-138, Siena, Italy, September 19-21, 2018.
9. Lukas Tuggener, Mohammadreza Amirian, Fernando Benites, Pius von Däniken, Prakhar Gupta, Frank-Peter Schilling, and Thilo Stadelmann. **"Design Patterns for Resource-Constrained Automated Deep-Learning Methods"**. **AI** section "Intelligent Systems: Theory and Applications" 1(4):510-538, MDPI, Basel, Switzerland, Novemer 06, 2020.
10. Dano Roost, Ralph Meier, Giovanni Toffetti Carughi, and Thilo Stadelmann. **"Combining Reinforcement Learning with Supervised Deep Learning for Neural Active Scene Understanding"**. In: Proceedings of the Active Vision and Perception in Human(-Robot) Collaboration Workshop at IEEE RO-MAN 2020 (**AVHRC'20**), online, August 31, 2020.

DIR3CT: Deep Image Reconstruction through X-Ray Projection-based 3D Learning of Computed Tomography Volumes

Collaboration with Inst. of Appl. Math. & Physics

- Topic: Compensation of motion artefacts in 3D CBCT reconstructed volumes using deep learning
- InnoSuisse, total volume **1.13 MCHF**
- Duration: 02/2020 – 05/2022
- Industry partner **Varian Medical Systems** (world market leader radiation therapy)
- Two involved ZHAW institutes **CAI & IAMP** (approx. 8 ZHAW researchers involved)
  - Focus CAI: 3D reconstruction using deep learning (supervised & unsupervised)
  - Focus IAMP: Physical modeling and simulation of motion, anatomical constraints
- Highly ambitious and technologically challenging

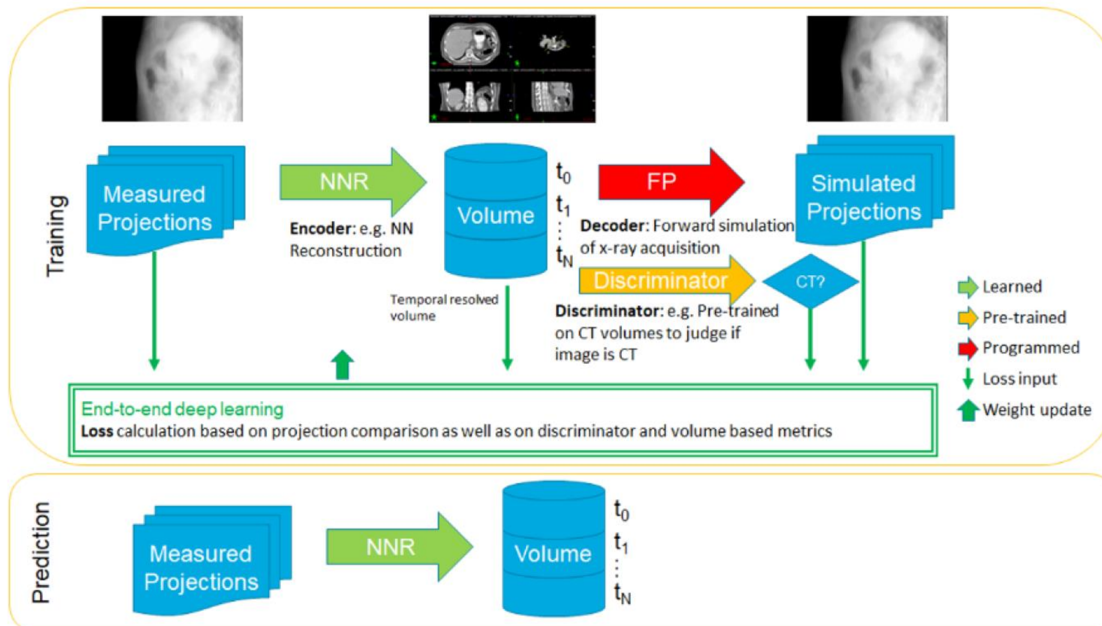

Zürcher Hochschule
für Angewandte Wissenschaften

Food Waste Analysis
Collaboration with the Inst. Of Embedded Systems

- Automatic detection of food waste in restaurants
- Embedded Machine Learning for waste classification
- Savings potential: At least CHF 2,500 per month per kitchen
- Research: Embedded System Design with GPU Edge Processing, Automatic Food Waste Classification
- Joint Innosuisse Project InES / CAI, July 19 – Aug 21

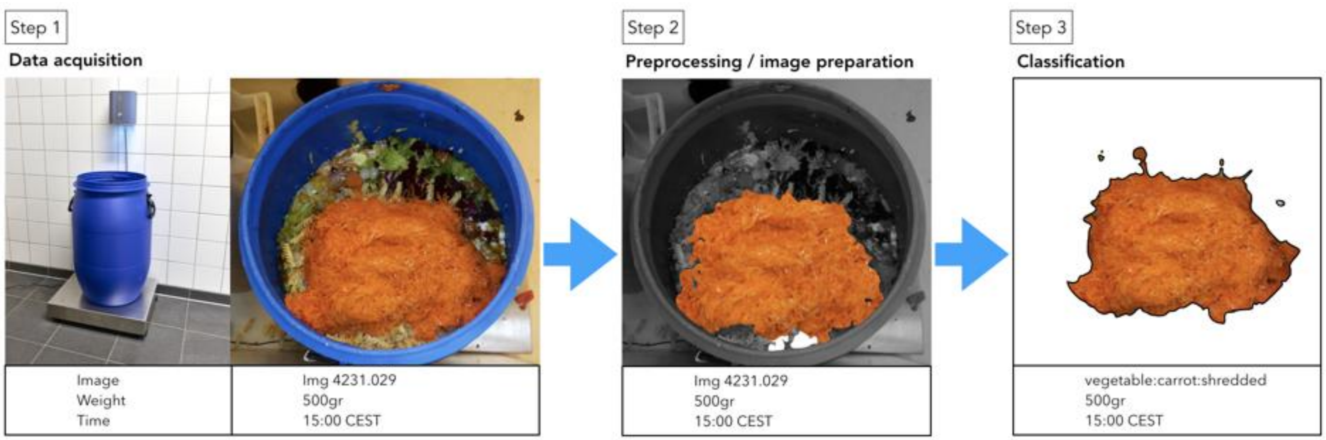

KITRO

www.kitro.ch

ADA: Automated Data Analyst
Collaboration with EPFL MLO Lab

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
- Result: top ranks in Google AutoDL'2020 competition

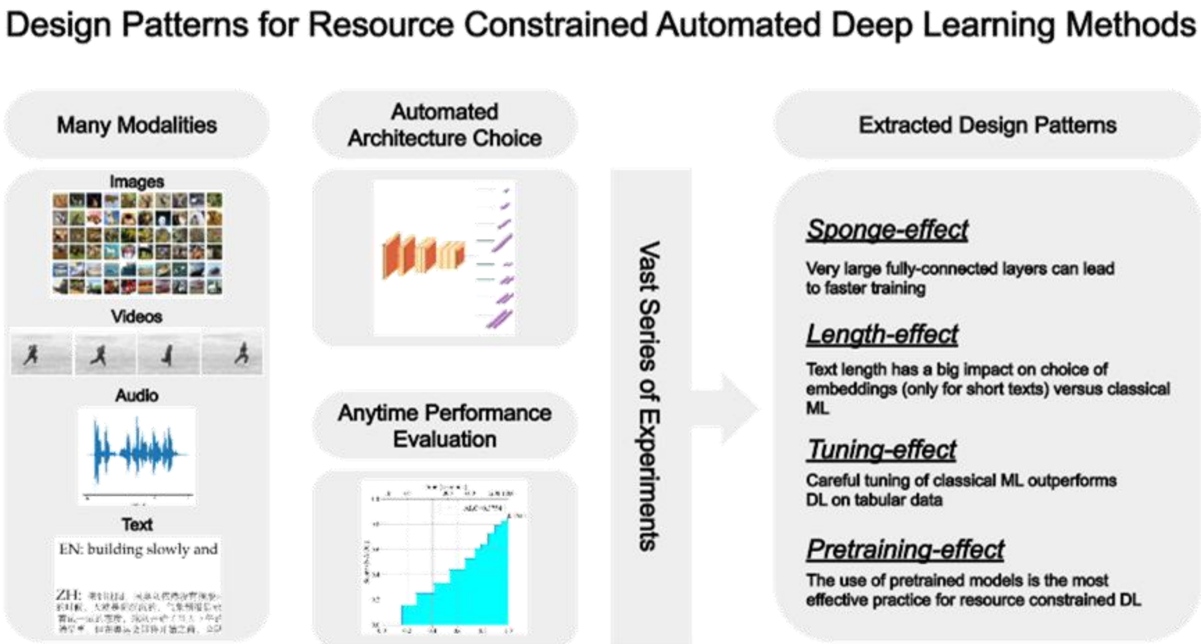

Stadelmann, Amirian, Arabaci, Arnold, Duivesteijn, Elezi, Geiger, Lörwald, Meier, Rombach & Tuggener (2018). *«Deep Learning in the Wild»*. ANNPR'2018.
Tuggener, Amirian, Rombach, Lörwald, Varlet, Westermann & Stadelmann (2019). *«Automated Machine Learning in Practice: State of the Art and Recent Results»*. SDS'19.
Tuggener, Amirian, Benites, von Däniken, Gupta, Schilling & Stadelmann (2020). *«Design Patterns for Resource-Constrained Automated Deep-Learning Methods»*. AI 1(4).

Zürcher Hochschule
für Angewandte Wissenschaften

LIBRA: Face matching & anti-spoofing
Collaboration with Inst. of Appl. Math. & Physics

[!] DEEPIMPACT

Schweizerische Eidgenossenschaft
Confédération suisse
Confederazione Svizzera
Confederaziun svizra

Swiss Confederation

Innosuisse – Swiss Innovation Agency

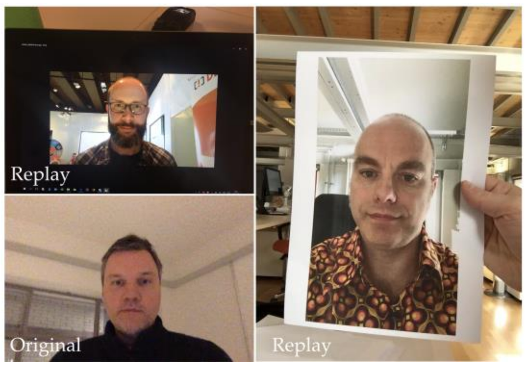

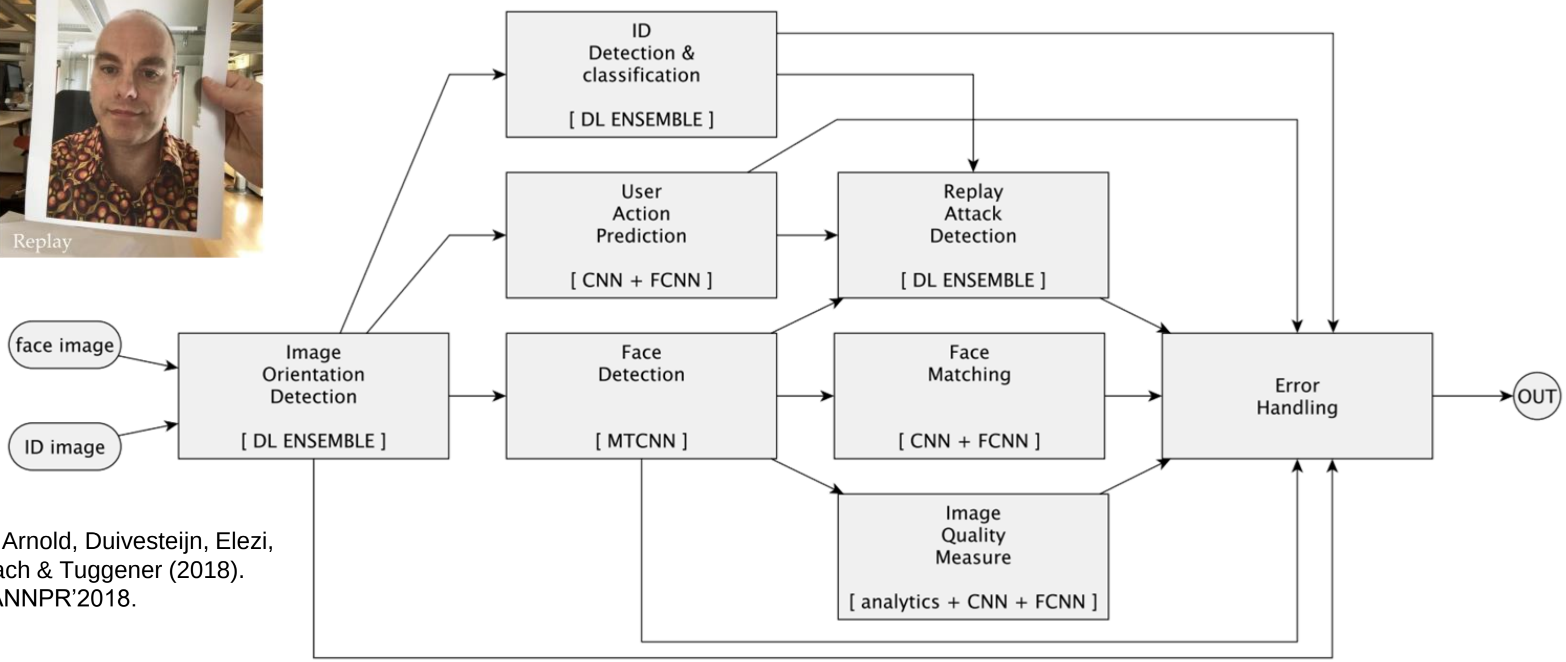

Stadelmann, Amirian, Arabaci, Arnold, Duivesteijn, Elezi, Geiger, Lörwald, Meier, Rombach & Tuggener (2018). *«Deep Learning in the Wild»*. ANNPR’2018.

DeepScore – Music OCR via Deep Neural Nets
Collaboration with IDSIA

Goal: Raise the accuracy of optical music recognition (OMR) by one order of magnitude to facilitate paper-free work of professional musicians

Challenge: Transfer the recent success of deep learning methods on numerous pattern recognition tasks (e.g., OCR) to the domain of music notation (which is 2D, without benchmarks, many syntactical constraints)

Solution: Enhance the open music scanner Audiveris by a new symbol classifier and segmenter based on convolutional neural networks to output musicXML

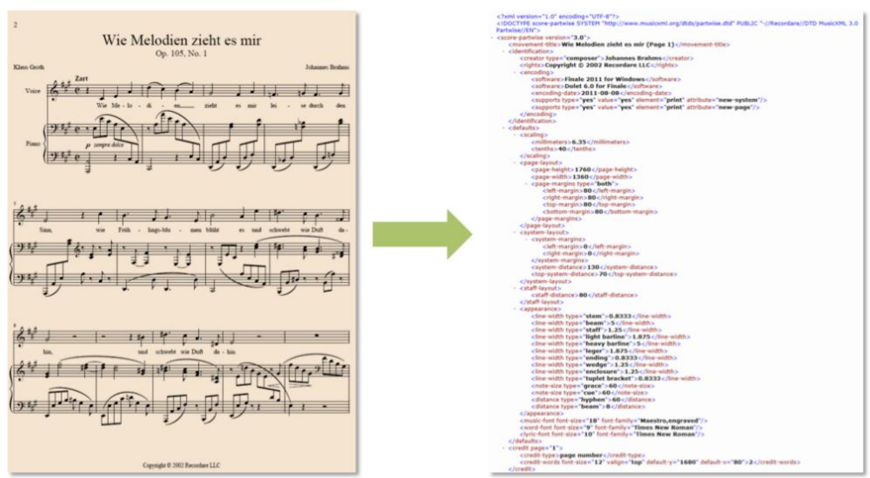

DeepScore – challenges & solutions

,

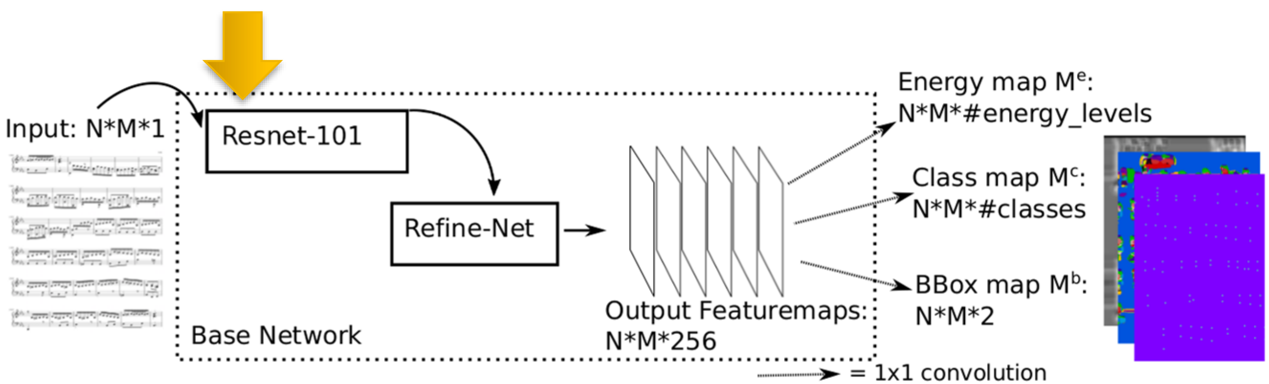

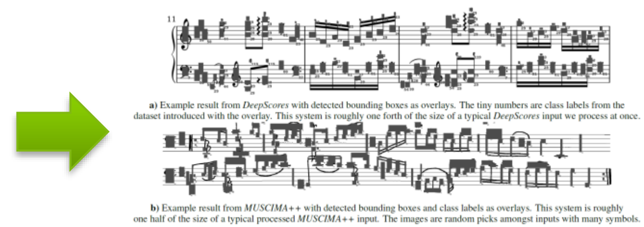

Tuggener, Elezi, Schmidhuber, Pelillo & Stadelmann (2018). *«DeepScores – A Dataset for Segmentation, Detection and Classification of Tiny Objects»*. ICPR'2018.
Tuggener, Elezi, Schmidhuber & Stadelmann (2018). *«Deep Watershed Detector for Music Object Recognition»*. ISMIR'2018.
Tuggener, Satyawan, Pacha, Schmidhuber & Stadelmann (2020). *«The DeepScoresV2 Dataset and Benchmark for Music Object Detection»*. ICPR'2020.

Zürcher Hochschule
für Angewandte Wissenschaften

SCAI: Smart Contract Analytics using Artificial Intelligence

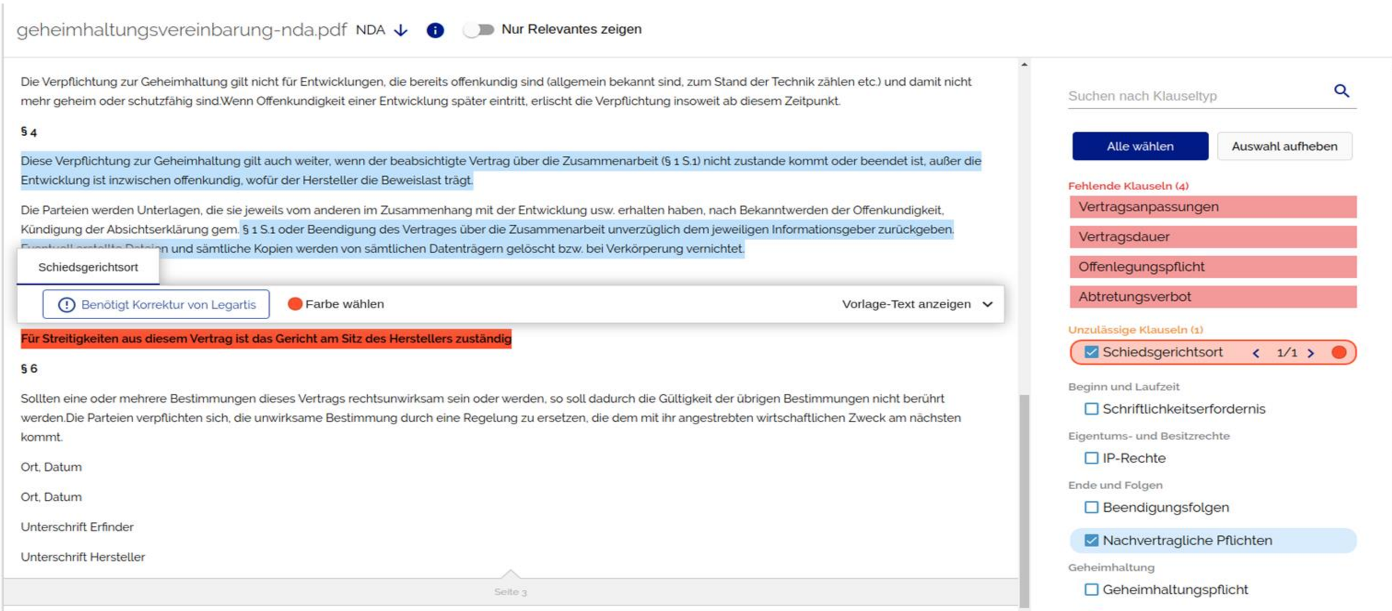

Innosuisse project (480'000 CHF)

Multi-label text classification: Classify contractual provisions (120+ labels)
Outlier detection: Find problematic provisions
Entity recognition: Detect companies, costs, penalties, jurisdiction etc.
Multilingual: EN, DE

LIHLITH: Lifelong Learning for Dialogue Systems

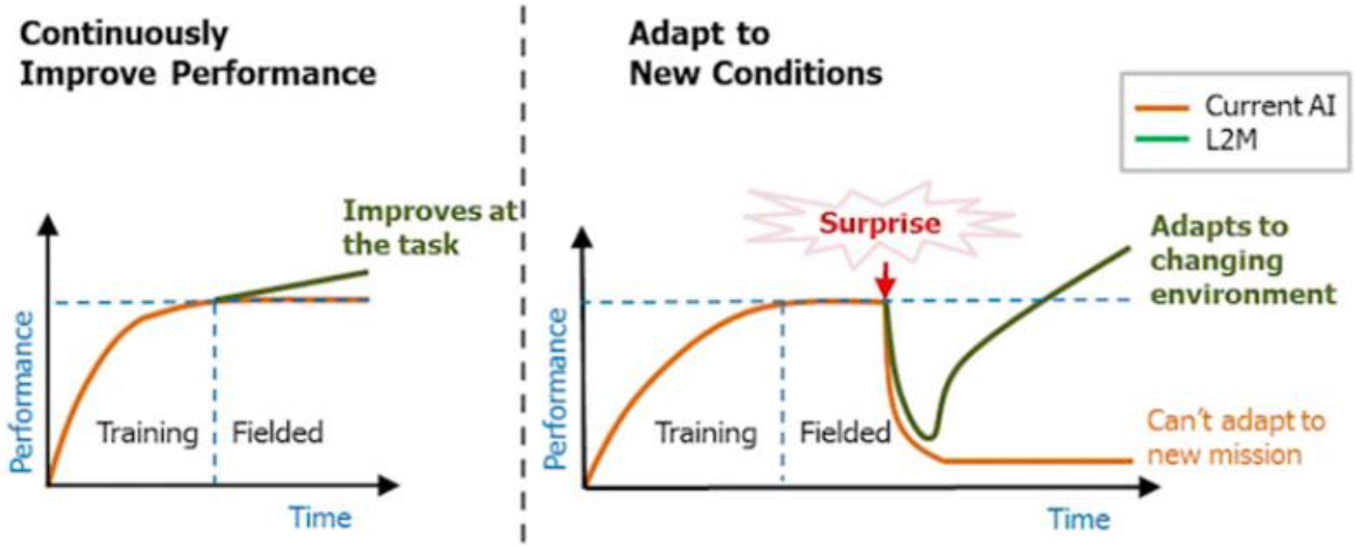

EU CHIST-ERA and SNF project (220'000 CHF)
- Fundamental research project
- What happens with dialogue systems after deployment? How does it learn new things continuously and autonomously?
- How to react when the algorithm is confronted with an unknown situation?
- Our contribution: benchmark to evaluate lifelong machine learning for natural language interfaces to databases

QualitAI
Optical Quality Control for MedTech Products

Goal: semi-automatic quality control of industrial goods with computer vision
Challenge: Work with small amounts of imbalanced data

Approach:
- Use state-of-the art deep learning models
- Use transfer learning, few-shot learning, image improvement to enable small data app

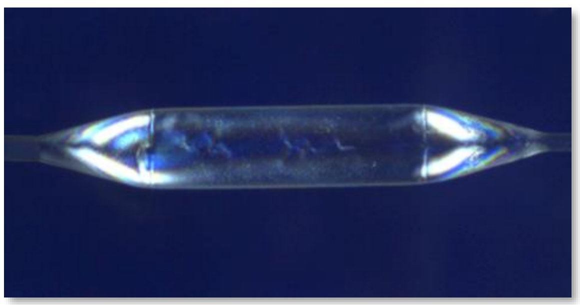

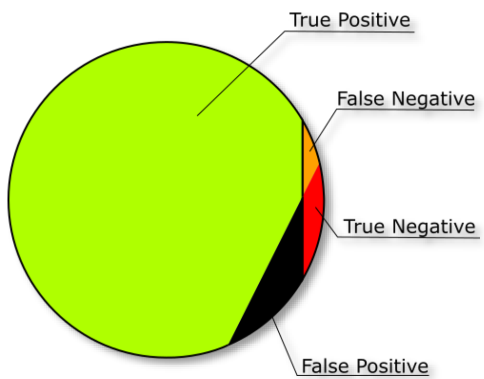

QualitAI – enabling model interpretability

- 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

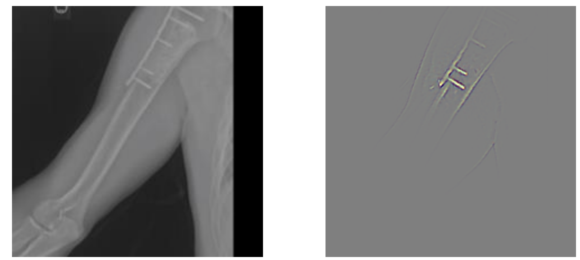

- Defends against adversarial attacks
  → thresholding local spatial entropy easily detects many adversarial attacking schemes through «lost focus»

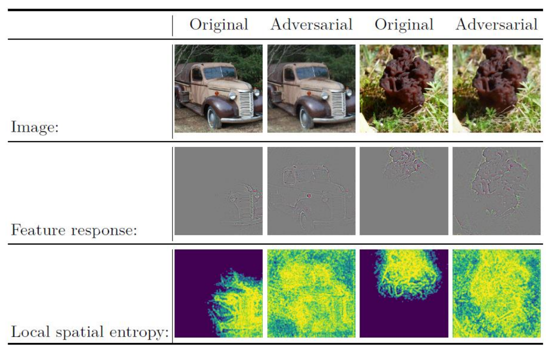

Stadelmann, Amirian, Arabaci, Arnold, Duivesteijn, Elezi, Geiger, Lörwald, Meier, Rombach & Tuggener (2018). *«Deep Learning in the Wild»*. ANNPR'2018.
Amirian, Schwenker & Stadelmann (2018). *«Trace and Detect Adversarial Attacks on CNNs using Feature Response Maps»*. ANNPR'2018.
Amirian, Tuggener, Chavarriaga, Satyawan, Schilling, Schwenker, & Stadelmann (2021). «Two to Trust: AutoML for Safe Modelling and Interpretable Deep Learning for Robustness». ECAI'2020 workshops.

Zürcher Hochschule
für Angewandte Wissenschaften

FarmAI: Automatic game playing
Collaboration with Inst. for Data Analysis & Process Design

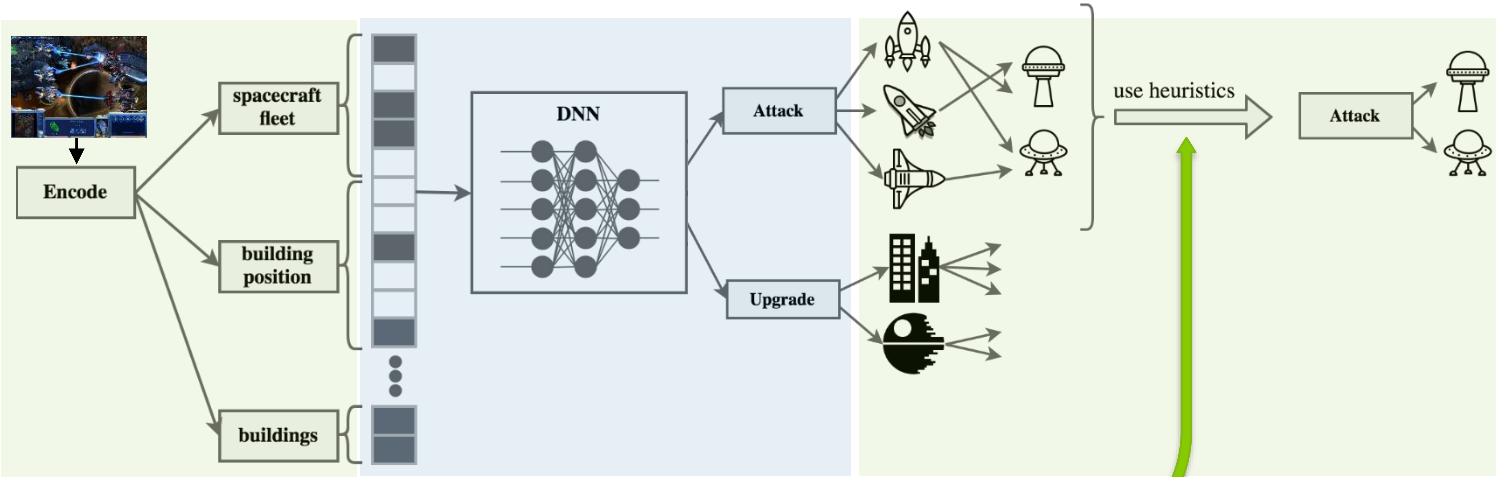

Reinforcement learning: deep Q network

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.

Project example: PANOPTES
Newspaper article segmentation for print media monitoring

Goal
- **Automatically segment newspaper pages** into constituting articles for automatic print media monitoring

Approach
- **Image-based** approach with **deep neural networks** that learn layouting principles from examples

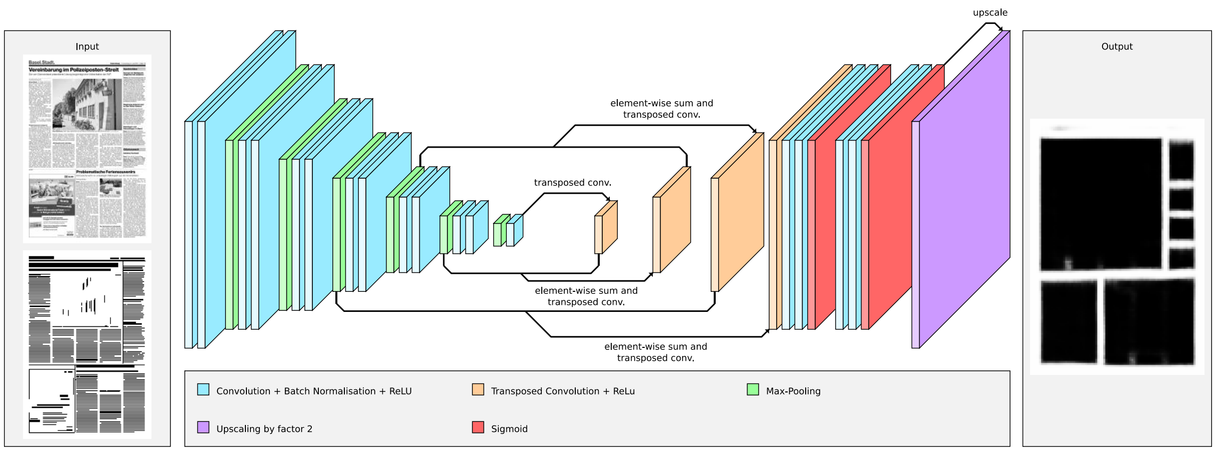

Meier, Stadelmann, Stampfli, Arnold, & Cieliebak. *"Fully convolutional neural networks for newspaper article segmentation"*. ICDAR 2017.

Stadelmann, Tolkachev, Sick, Stampfli, & Dürr. *"Beyond ImageNet - Deep Learning in Industrial Practice"*. In: Braschler et al. (Eds). "Applied Data Science – Lessons Learned for the Data-Driven Business", Springer, 2019.

Zürcher Hochschule
für Angewandte Wissenschaften

Project example: Complexity 4.0
Collaboration with HSG et al.

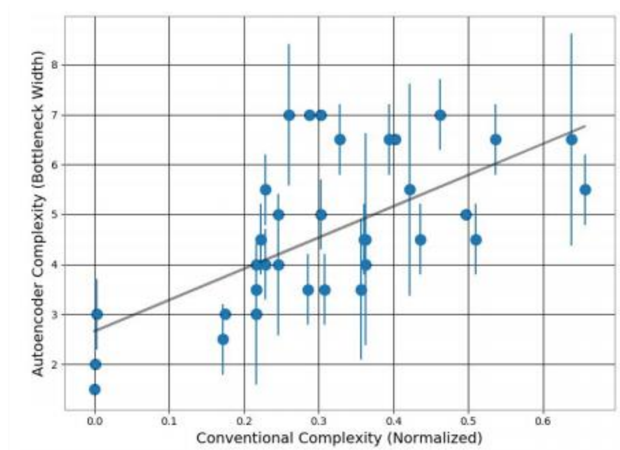

Schweizerische Eidgenossenschaft
Confédération suisse
Confederazione Svizzera
Confederaziun svizra

Swiss Confederation

Innosuisse – Swiss Innovation Agency

University of St.Gallen

datalab
www.zhaw.ch/datalab

data service
Swiss Alliance for Data-Intensive Services

Goal

- **Reduce** unnecessary **complexity of product variability** in production environments in a data-driven (~automatable) fashion

Project team

- Business partners: **2 different industries** with large production facilities in CH
- **Economists**: ITEM-HSG (technology management, business models)
- **Engineers**: ZHAW-Engineering (machine learning), ZHAW-Life Sciences (simulation)

Results

- *“The paradigm of **data-driven decision support** can […] enter the domain of a highly qualified business consultant, **deliver**ing the **quantitative results** necessary to ponder informed **management decisions**.”*
- *“**It is merely the knowledge** of what methods and technologies are possible and available **that** currently **hinders the faster adoption** of the data-driven paradigm in businesses.”*

Hollenstein, Lichtensteiger, Stadelmann, Amirian, Budde, Meierhofer, Füchslin, & Friedli *“Unsupervised Learning and Simulation for Complexity Management in Business Operations”*. In: Braschler et al. (Eds). “Applied Data Science – Lessons Learned for the Data-Driven Business”, Springer, 2019.

Talkalyzer
Contact: Prof. Dr. Thilo Stadelmann

Goal: Speaker Recognition in meetings on mobile devices
Challenge: Build reliable speaker models

Approach:
- Loosen iid. assumption on feature vectors
- Use Deep Neural Network approach on continuous audio features
  ➔ find typical sounds of a speaker in a spectrogram

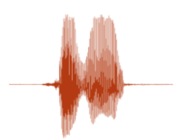
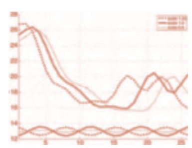
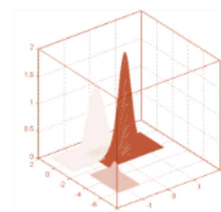
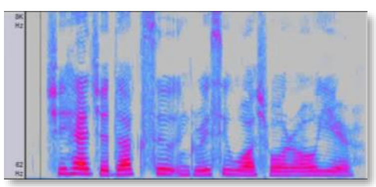
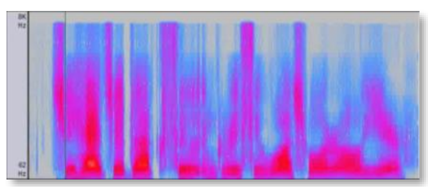
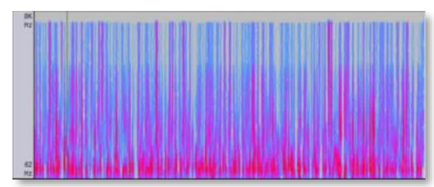

Talkalyzer – exploiting time information

CNN (MLSP'16)

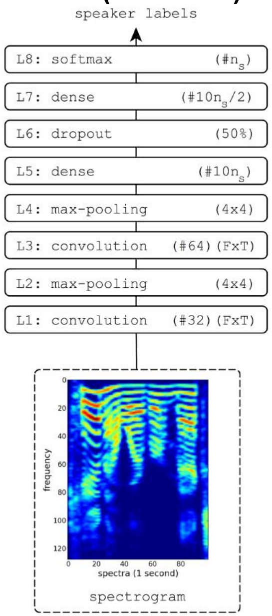

CNN & clustering-loss (MLSP'17)

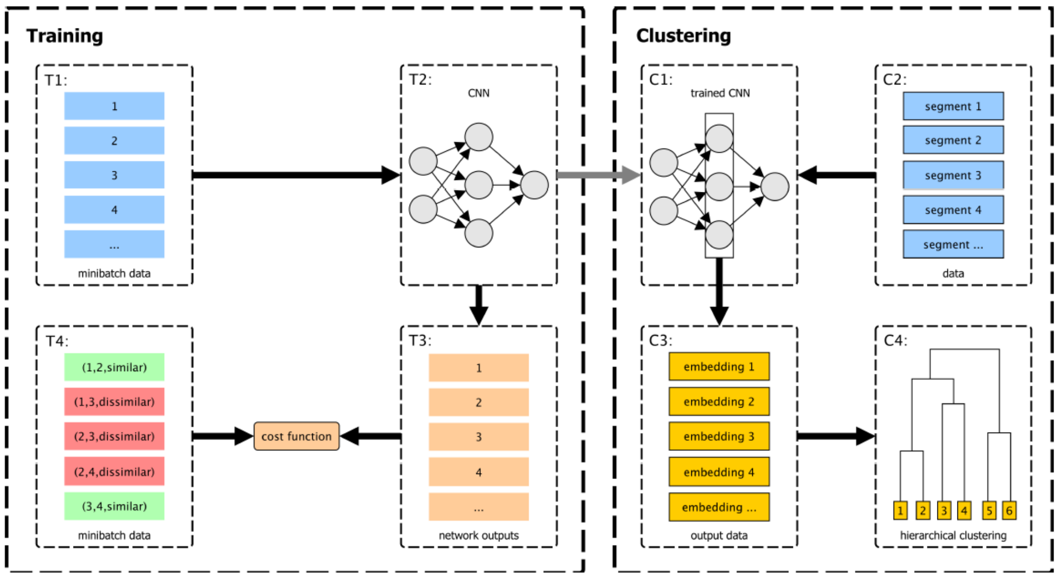

RNN & clustering-loss (ANNPR'18)

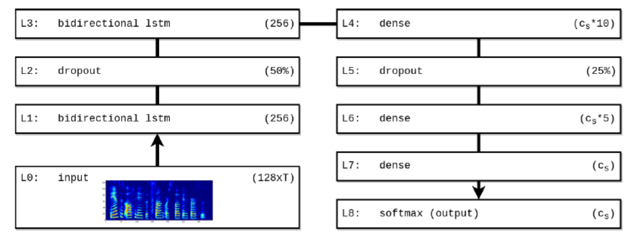

Method	MR	MR (legacy)
RNN /w PKLD	**2.19%** $\left(\frac{1.25\%+2.5\%+1.25\%+3.75\%}{4}\right)$	4.38% (average of 4 runs)
CNN /w PKLD [24]	-	**5%**
CNN /w cross entropy [23]	-	**5%**
ν-SVM [40]	6.25%	-
GMM/MFCC [40]	12.5%	-

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.

Zürcher Hochschule
für Angewandte Wissenschaften

Sentiment Analysis

Contact: Dr. Mark Cieliebak

Challenge:

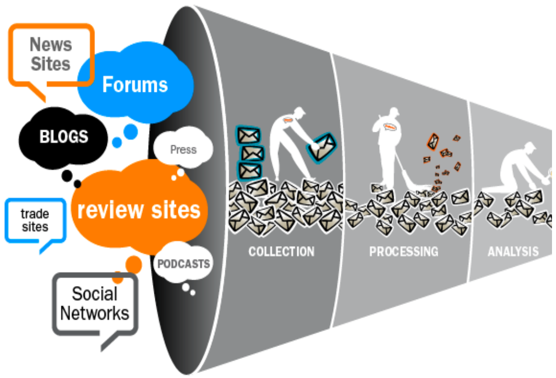

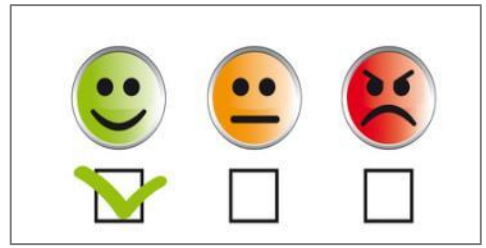

Approach:

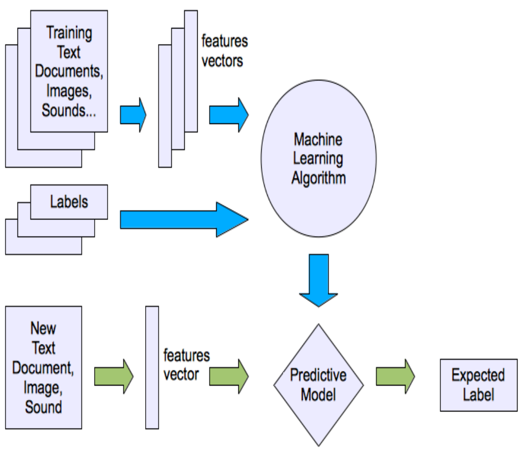

Demo:

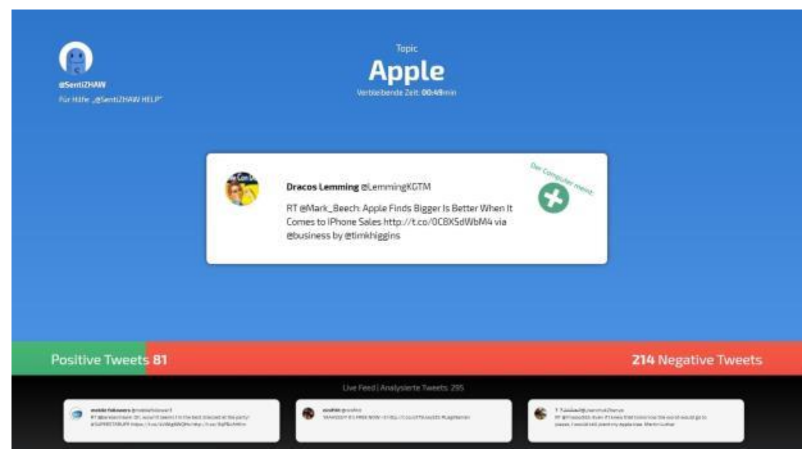

SODES – Swiss Open Data Exploration System

Contact: Prof. Dr. Mark Cieliebak

Challenge: Open Data promises to be a gold mine – but accessing and combining data from different data sources turns out to be non-trivial and very time consuming

Goal: A platform that enables easy and intuitive access, integration and exploration of different data sources

Solution:

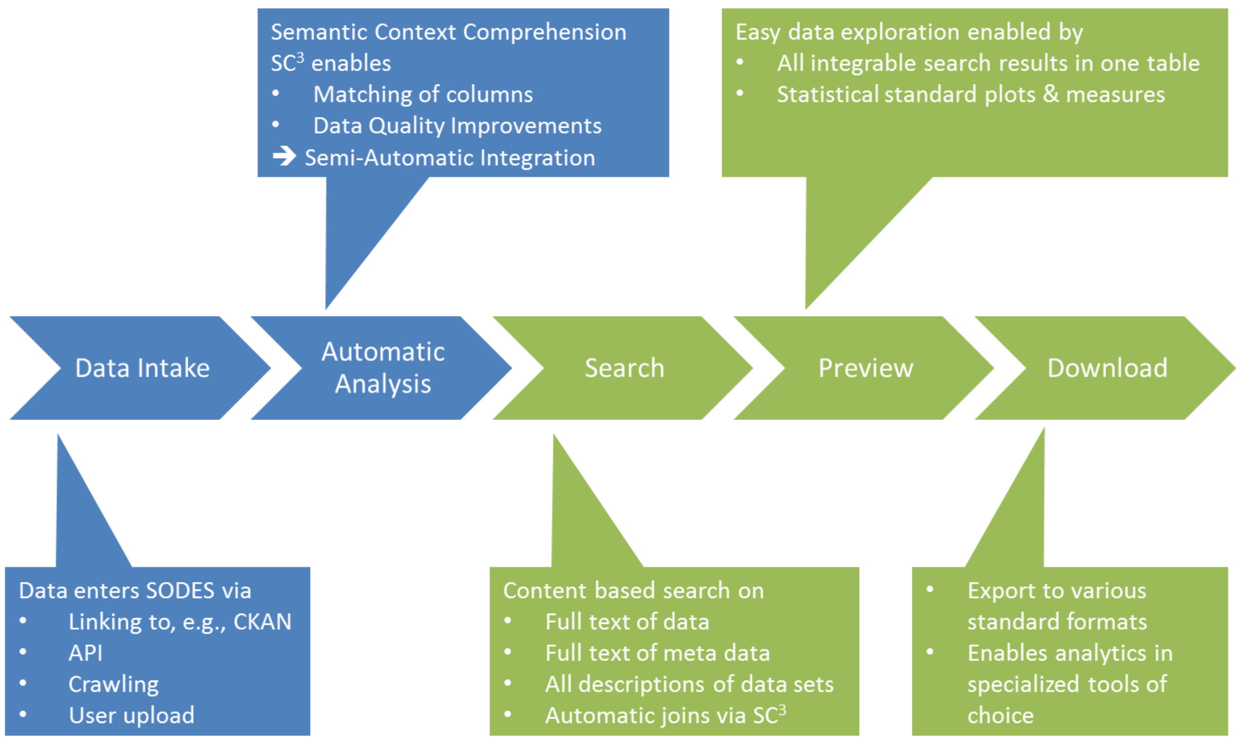

Zürcher Hochschule
für Angewandte Wissenschaften

DaCoMo – Data-driven Condition Monitoring
Contact: Prof. Dr. Thilo Stadelmann

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

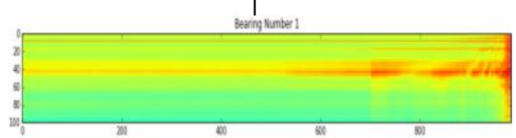

e.g., RNN autoencoder

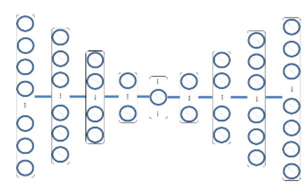

early detection of fault

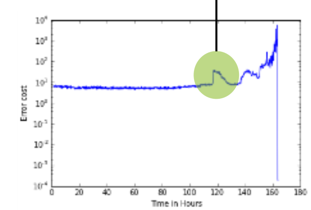

Stadelmann, Tolkachev, Sick, Stampfli, & Dürr. *"Beyond ImageNet - Deep Learning in Industrial Practice"*. In: Braschler et al. (Eds). "Applied Data Science – Lessons Learned for the Data-Driven Business", Springer, 2019.

Influencer Detection in Social Media

Target Specific, Interactive

Contact: Dr. Mark Cieliebak

Zürcher Hochschule
für Angewandte Wissenschaften

zh
aw

1 Person

1 Story

1 Blog Post

1 Tweet

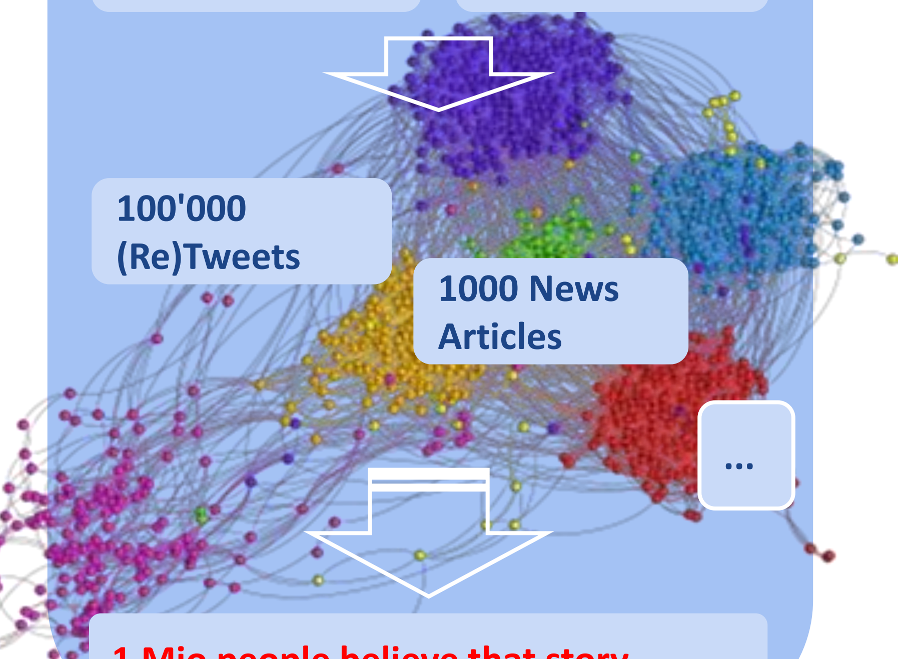

100'000
(Re)Tweets

1000 News
Articles

...

1 Mio people believe that story

How
to find
this
person
in
1 second

?

Social Media Monitoring

Data Source APIs

Periodic Data Fetch

Analyze Unstructured Text

NLP

Full-Text
Index

Graph
Structure

Fast Retrieval