Lessons Learned from Watching Machines Learn

AMLD EPFL 2022, AI & Mobility Track, Lausanne, Switzerland, March 29, 2022

Thilo Stadelmann

Zürcher Hochschule
für Angewandte Wissenschaften

Zürcher Hochschule
für Angewandte Wissenschaften

Agenda

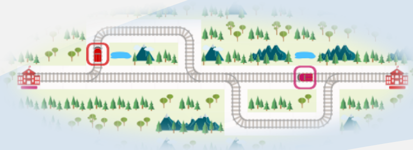

2. Communication
in multi-agent RL

3. Discussion

1. Computer vision
architecture design

We created a number of practical deep learning applications over the years…

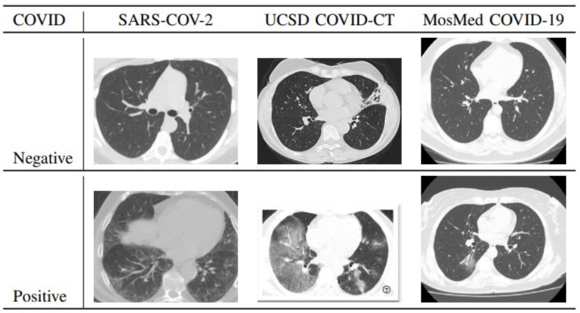

Medical imaging:
domain adaptation for diagnosis

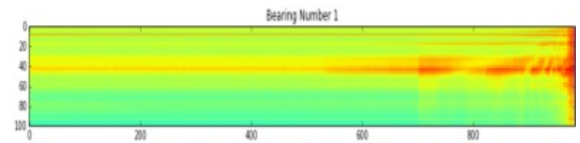

Industrial vision: data-driven predictive maintenance

Industrial vision:
defect recognition

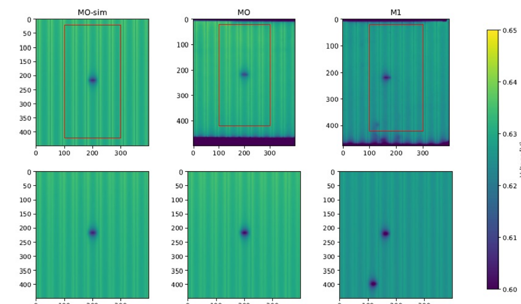

Industrial vision: prediction of solar cell simulation parameters from a real-world picture

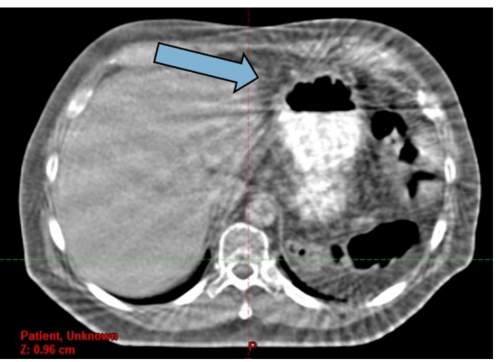

Medical imaging:
motion artifact reduction

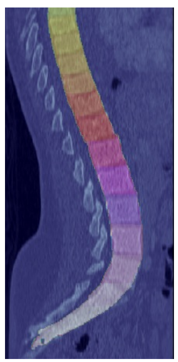

Medical imaging:
vertebrae detection

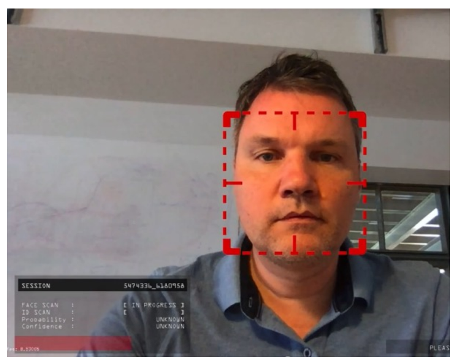

Biometrics:
robust face recognition

Industrial vision: food waste segmentation

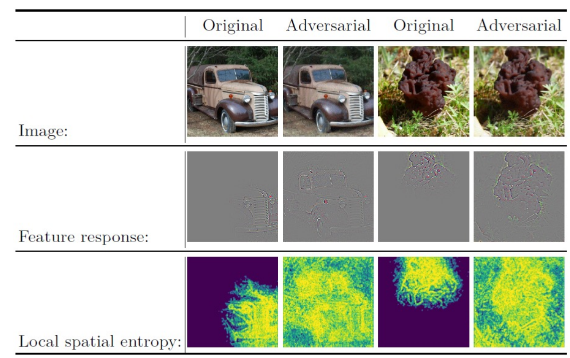

Industrial vision: explainability and adversarial attack detection

Is ImageNet a good basis for deriving CNN architectures for other use cases?

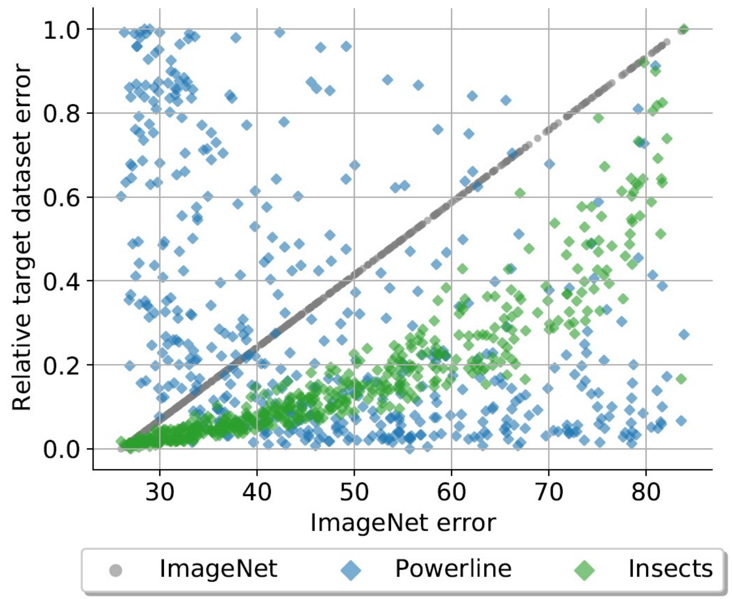

Fig. 1. Is a CNN *architecture* that performs well on ImageNet automatically a good choice for a different vision dataset? This plot suggests otherwise: It displays the relative test errors of 500 randomly sampled CNN architectures on three datasets (ImageNet, Powerline, and Insects) plotted against the test error of the same architectures on ImageNet. The architectures have been trained from scratch on all three datasets. Architectures with low errors on ImageNet also perform well on Insects, on Powerline the opposite is the case.

Tuggener, Schmidhuber & Stadelmann: *„ImageNet as a Representative Basis for Deriving Generally Effective CNN Architectures"*, under review, 2022.

Is ImageNet… (contd.)
Study design and results

- 500 randomly sampled architectures from the AnyNetX family (incl. AlexNets, VGGs, ResNets, RegNets)

- Trained from scratch on ImageNet and 8 relevant real-world datasets

DATASET	NO. IMAGES	NO. CLASSES	IMG. SIZE
CONCRETE	40K	2	227×227
MLC2008	43K	9	312×312
IMAGENET	1.3M	1000	256×256
HAM10000	10K	7	296×296
POWERLINE	8K	2	128×128
INSECTS	63K	291	296×296
NATURAL	25K	6	150×150
CIFAR10	60K	10	32×32
CIFAR100	60K	100	32×32

- Tested on (a) a test set from ImageNet and (b) on the same type used for training

- Extensive ablation studies to show validity

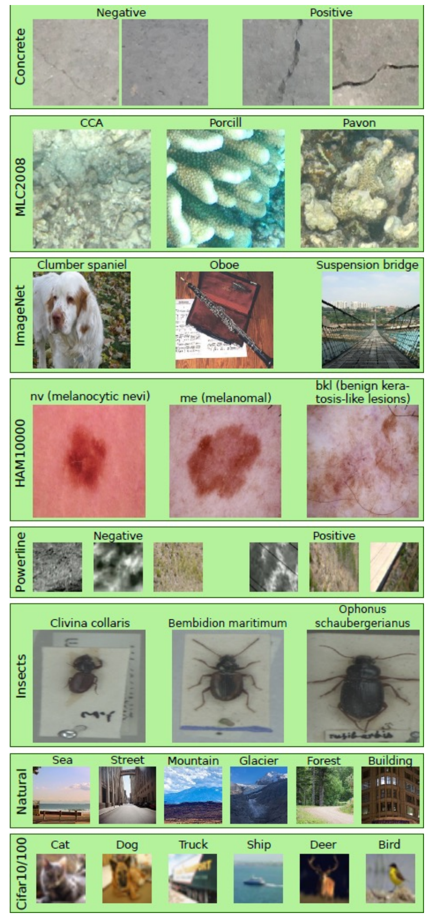

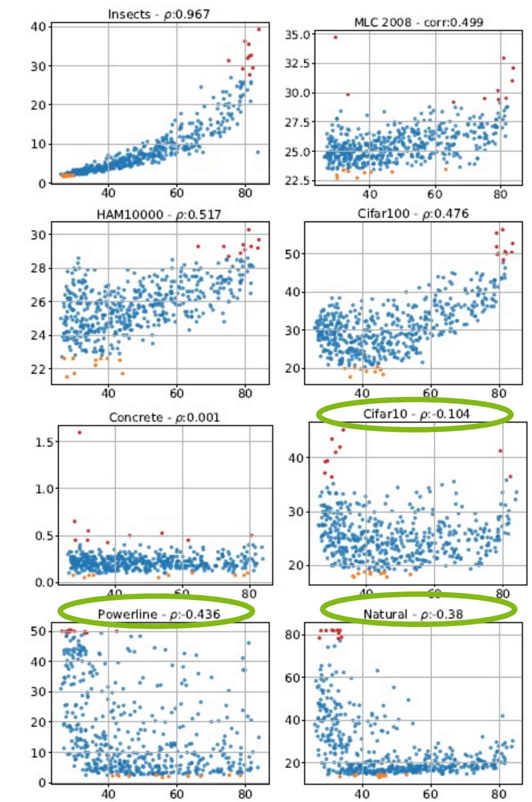

Fig. 4. Test errors of all 500 sampled architectures on target datasets (y-axis) plotted against the test errors of the same architectures (trained and tested) on ImageNet (x-axis). The top 10 performances on the target datasets are plotted in orange and the worst 10 performances in red.

Is ImageNet… (contd.)
Findings

- **Architecture search based on ImageNet** performance **is worse than random search** for at least Natural, Powerline and Cifar10

- **Varying the number of classes in ImageNet** is a cheap and **effective remedy** (i.e., randomly selecting x classes and deleting the rest of the dataset → ImageNet-x)
- …whereas **image-similarity or image size** play **not** an **important** role (e.g., Natural images are most sikmilar to ImageNet's)

- **Hyperparameters cumulative block depth and** cumulative block **width** can **drastically change based on dataset** and are influenced by class count

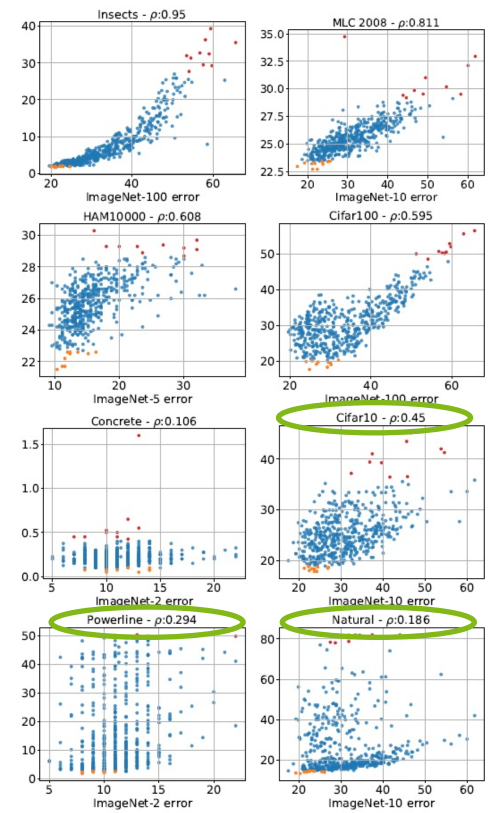

Fig. 6. Test errors of all 500 sampled architectures on target datasets (y-axis) plotted against the test errors of the same architectures on the ImageNet-X (x-axis). The top 10 performances on the target dataset are orange, the worst 10 performances red.

Agenda

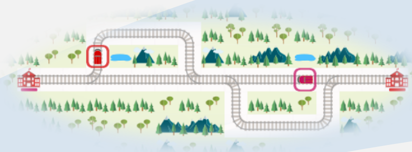

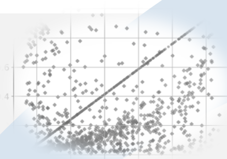

1. Computer vision architecture design

2. Communication in multi-agent RL

3. Discussion

Mutli-agent RL for train rescheduling

Problem description

- How to **adjust for small delays** ("rescheduling") automatically **in a** more and more **packed railway network** like the one of SBB?

- **Closed-form optimization impossible** due to combinatorial explosion of rerouting options

- RL still in its infancy for practical *high-consequence* environments → Flatland challenge to explore options

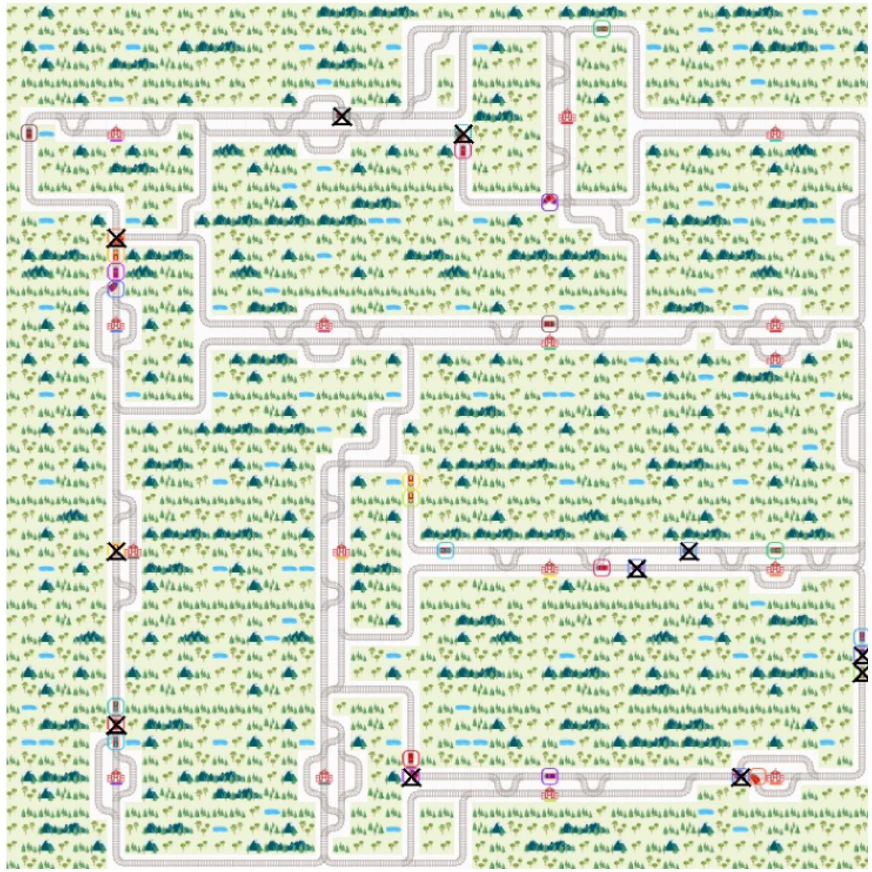

Lessons learned on RL in rescheduling
(based on a rank-6 entry to the Flatland challenge)

How to make RL sample-efficient:

- Using **task-specific heuristics** to present the agent with percepts only when a decision is necessary (i.e., at switches) increases the performance from 44.5% to 82.9%
- Using **curriculum learning** to learn fundamental behavior in easy environments and gradually increase complexity ensures rank 6/32 in the more realistic Flatland Round 2

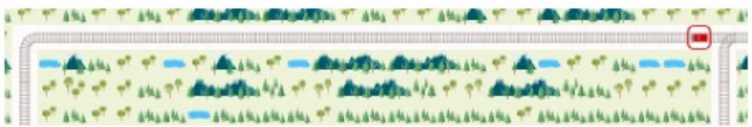

Screenshot from Flatland environment. A train heading to the left. The only reasonable action is to ride forward.

	Level 0	Level 1	Level 2	Level 3	Level 4
Next level on success rate	70%	70%	75%	70%	60%
Nr. of agent	4	8	12	16	20
Env. size	25x25	30x30	40x40	50x50	50x50
Num. cities	5	8	10	12	16
Max. rails between cities	1	2	2	2	2
Max. rails in city	2	2	3	3	3

General remark:

- **Policy gradient methods** seem **generally inappropriate** for high-consequence environments (i.e., one bad action leads to unresolvable catastrophes)
- Reason is **stochasticity**: if distributions over actions are learned and many agents are present in a single environment, the probability of having one bad action in every time step approaches certainty

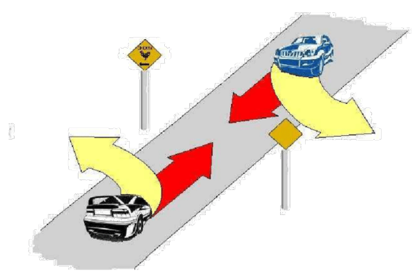

An emerging machine language?

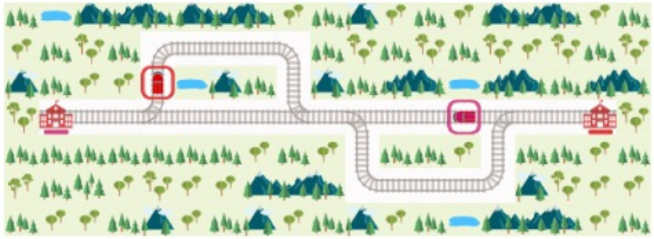

Humans would communicate to **negotiate** who would take **the detour**

What happens if we **add communication actions** (5 free tokens + EOT) and a shared **communication buffer in the observation** to the RL scenario**?**

Communication process:

1. **Communication loop** is entered upon first comm. action taken by any agent
2. Agents can sequentially **read** the comm. **buffer** and **add** a comm. **action**
3. Comm. loop **ends when** both agents issue the **EOT** action
4. **Then**, both agents can select regular (non-comm.) actions again and **proceed in the environment**

➔ **Does** the **general ability** to negotiate (i.e., exchange an arbitrary long sequence of tokens until mutually agreed to end) **help** in **practically** avoiding collision**?**

A first glimpse

Training

- Reward -1 if agents collide after negotiation; +1 otherwise
- Agents don't know who they are and need to take actions in parallel → cannot stick to go only one way or react to first mover
- 1M episodes training (A3C)

Results

- Success rate increases from 47% to 95%!
- High diversity in machine dialogues!
- (See examples on the right →)

Implications

- Allowing arbitrarily long sequences of 5 tokens can lead to a Turing-completeness
- But what happens actually?

Timestep	Actions agent 1|2	Outcome
0	4 | 2	
1	5 | 5	Success
0	3 | 0	
1	1 | 5	
2	5 | 5	Success
0	3 | 5	
1	5 | 5	Success
0	3 | 1	
1	3 | 2	
2	5 | 0	
3	5 | 5	Crash
0	3 | 2	
1	5 | 3	
2	5 | 4	
3	2 | 5	
4	5 | 5	Success
0	4 | 3	
1	3 | 1	
2	5 | 5	Success

Discussion

- What puzzling aspects of your research have you so far ignored in hunt of a different goal?
- Do you think there is a lesson to learn from searching for an explanation?
- Do you think it pays off to take these detours?

- Ideas for continuing the RL & communication work?

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

APPENDIX

Sample projects

Zürcher Hochschule
für Angewandte Wissenschaften

We created a number of practical deep learning applications over the years…

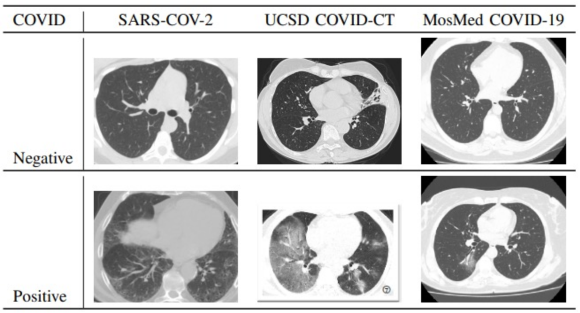

Medical imaging: domain adaptation for diagnosis

Document analysis: article segmentation

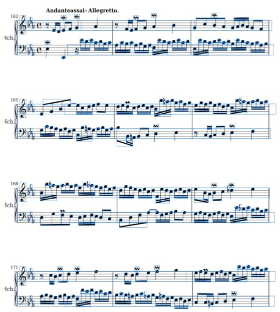

Document analysis: optical music recognition

Industrial vision: quality control

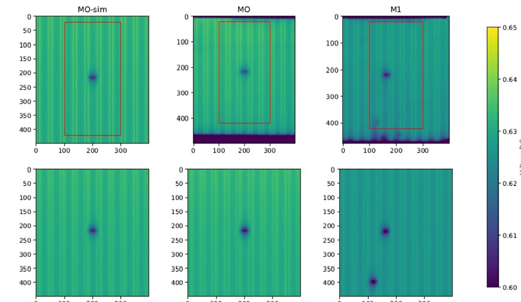

Industrial vision: prediction of solar cell simulation parameters from a real-world picture

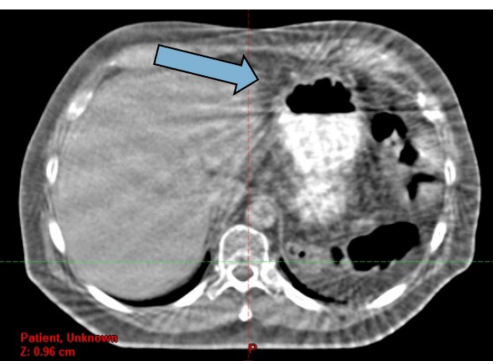

Medical imaging: motion artifact reduction

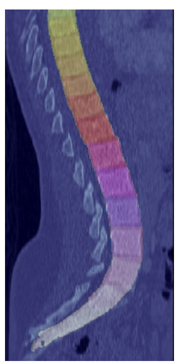

Medical imaging: vertebrae detection

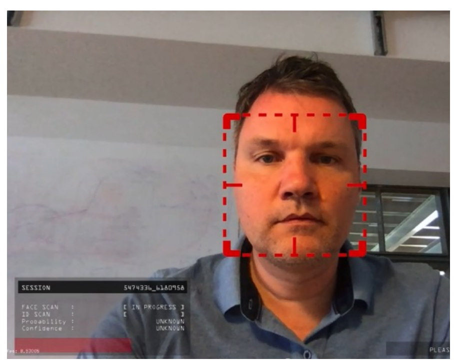

Biometrics: robust face recognition

Industrial vision: food waste segmentation

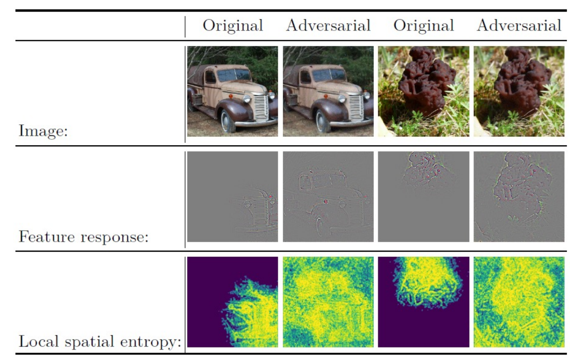

Industrial vision: explainability and adversarial attack detection

Zürcher Hochschule
für Angewandte Wissenschaften

We created a number of practical deep learning applications over the years… (contd.)

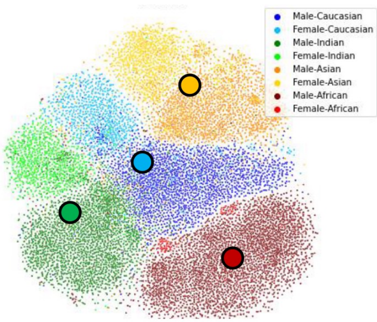

Biometrics: origins of bias in face recognition

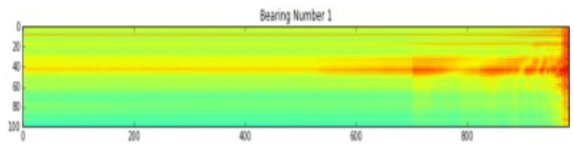

Industrial vision: data-driven predictive maintenance

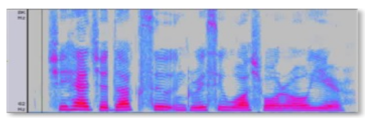

Biometrics: automatic speaker recognition

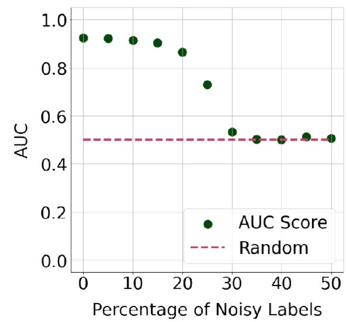

Industrial vision: learning with noisy labels

Document recognition: newspaper segmentation

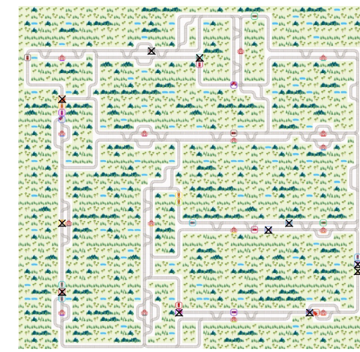

Logistics planning: train rescheduling

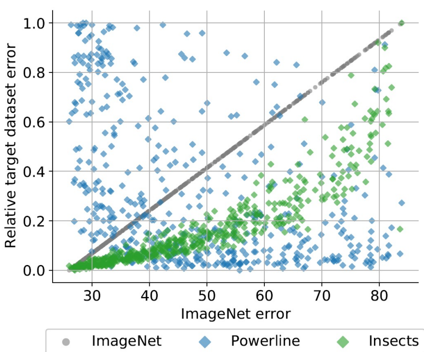

ML fundamentals: neural architecture design

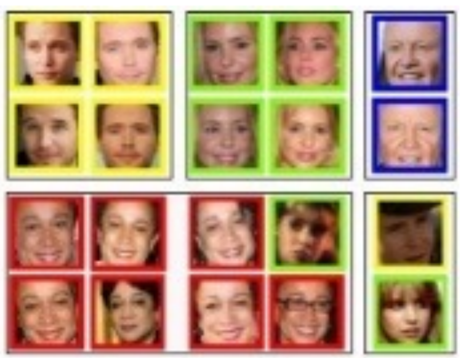

ML fundamentals: learning inductive biases for clustering

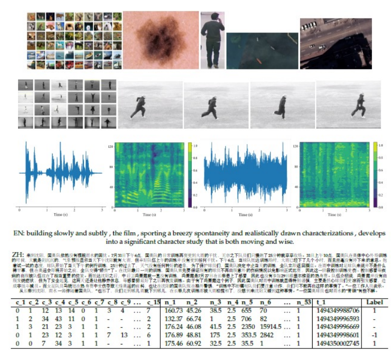

ML fundamentals: automated deep learning

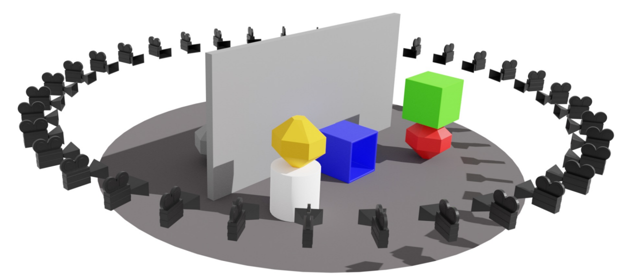

ML fundamentals: active learning for computer vision

Computer Vision, Perception & Cognition Group

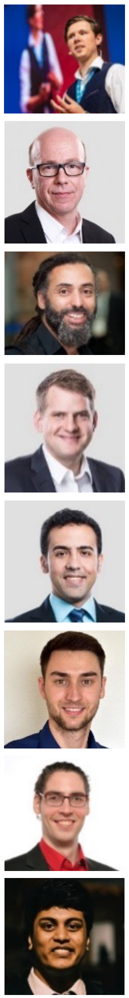

Machine learning-based Pattern Recognition

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

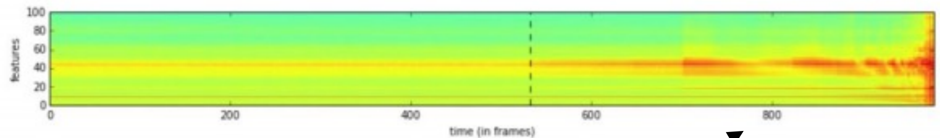

[3]

[1]

[2]

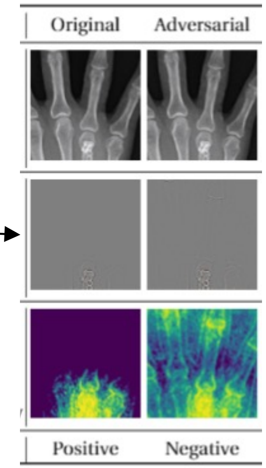

[5]

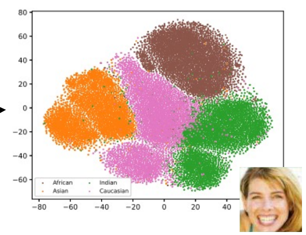

[4]

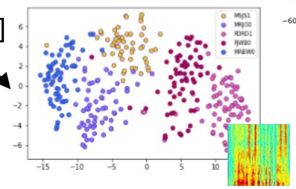

[7]

[6]

[1]

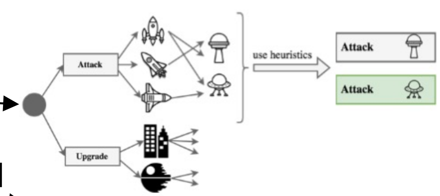

[10]

[8]

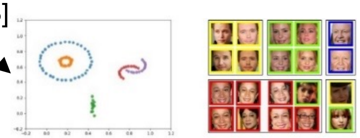

[9]

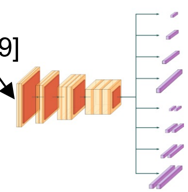

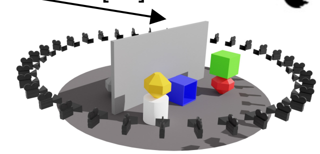

CVPC Group: references for overview

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.

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

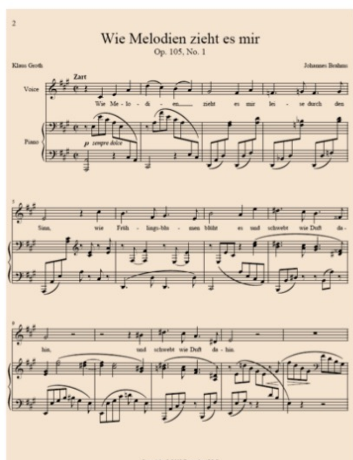

DeepScore – challenges & solutions

,

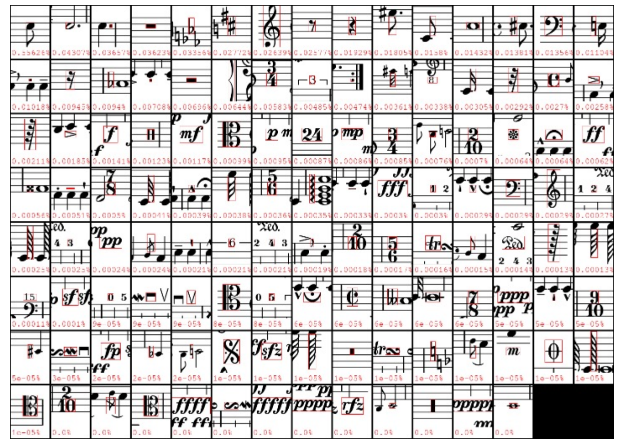

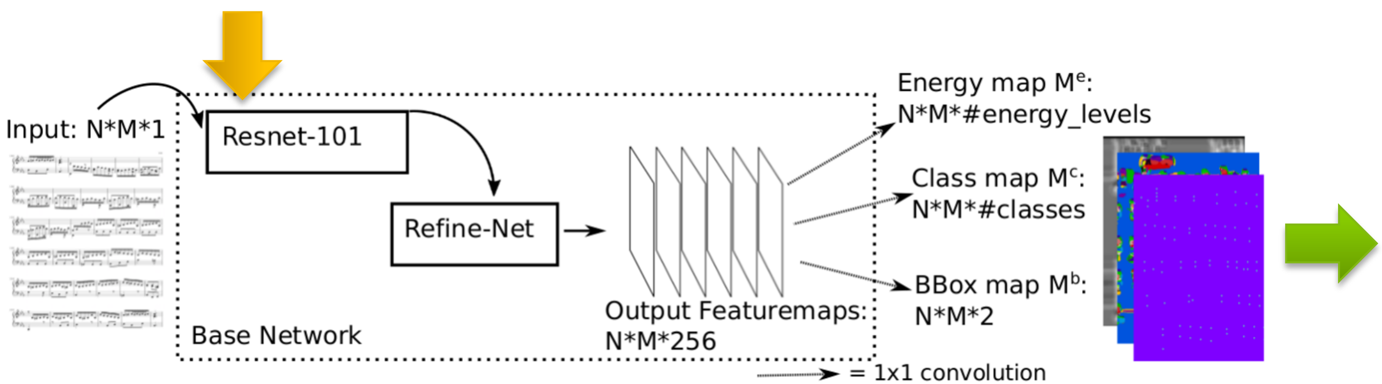

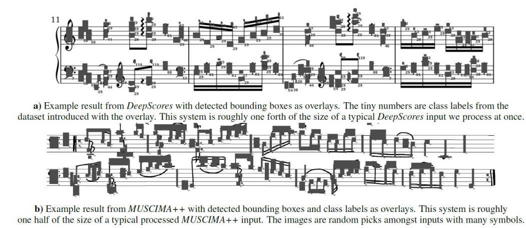

Tuggener, Elezi, Schmidhuber, Pelillo & Stadelmann (2018). *«DeepScores – A Dataset for Segmentation, Detection and Classification of Tiny Objects»*. ICPR’2018.
Tuggener, Elezi, Schmidhuber & Stadelmann (2018). *«Deep Watershed Detector for Music Object Recognition»*. ISMIR’2018.
Tuggener, Satyawan,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

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

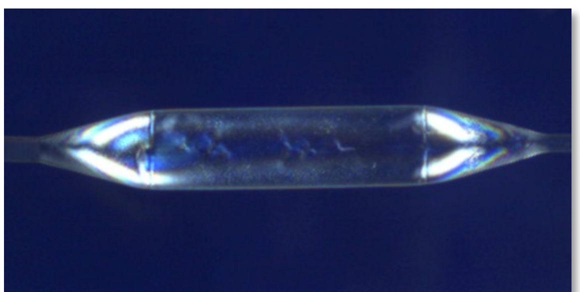

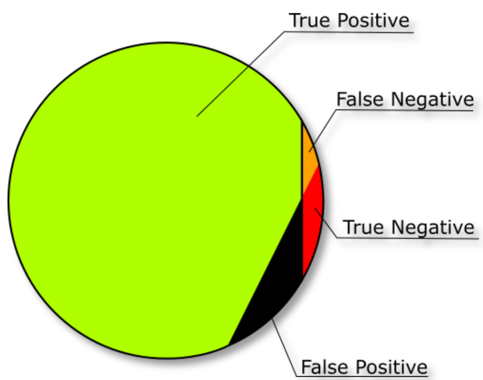

QualitAI – enabling model interpretability

- 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

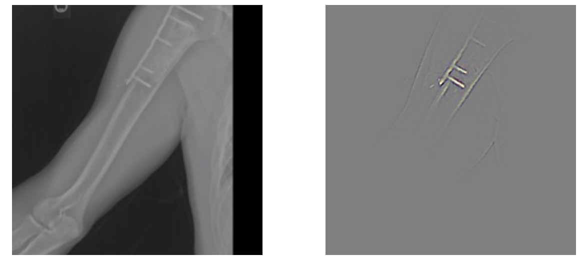

- Defends against adversarial attacks
  → thresholding local spatial entropy easily detects many adversarial attacking schemes through «lost focus»

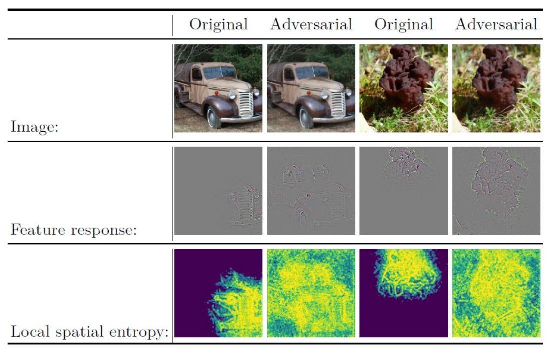

Stadelmann, Amirian, Arabaci, Arnold, Duivesteijn, Elezi, Geiger, Lörwald, Meier, Rombach & Tuggener (2018). *«Deep Learning in the Wild»*. ANNPR'2018.
Amirian, Schwenker & Stadelmann (2018). *«Trace and Detect Adversarial Attacks on CNNs using Feature Response Maps»*. ANNPR'2018.
Amirian, Tuggener,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

FarmAI: Automatic game playing

Collaboration with Inst. for Data Analysis & Process Design

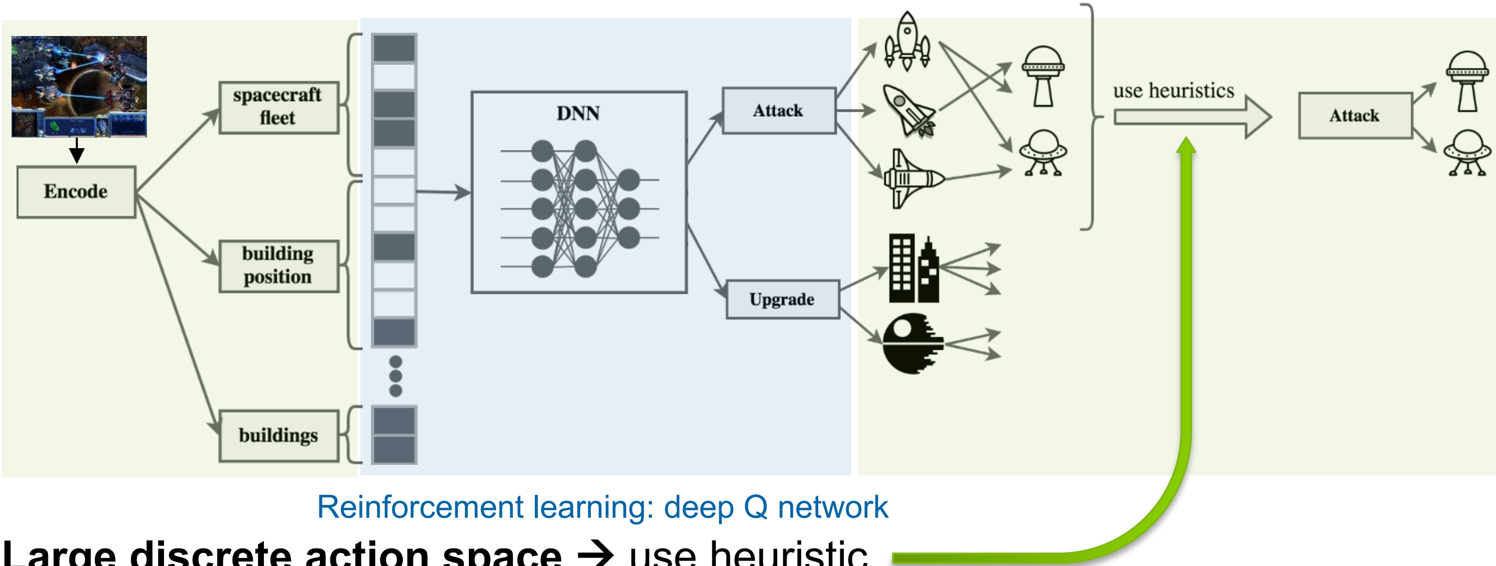

Reinforcement learning: deep Q network

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.

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

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

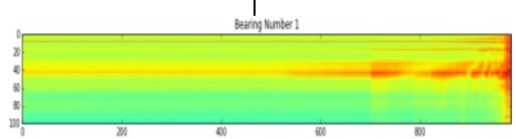

e.g., RNN autoencoder

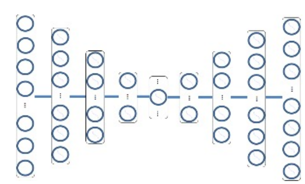

early detection of fault

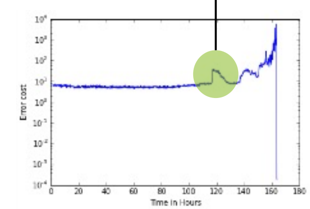

Stadelmann, Tolkachev, Sick, Stampfli, & Dürr. *"Beyond ImageNet - Deep Learning in Industrial Practice"*. In: Braschler et al. (Eds). "Applied Data Science – Lessons Learned for the Data-Driven Business", Springer, 2019.