Overcoming the Reproducibility Crisis in Sciences using AI?

Kristian Kersting

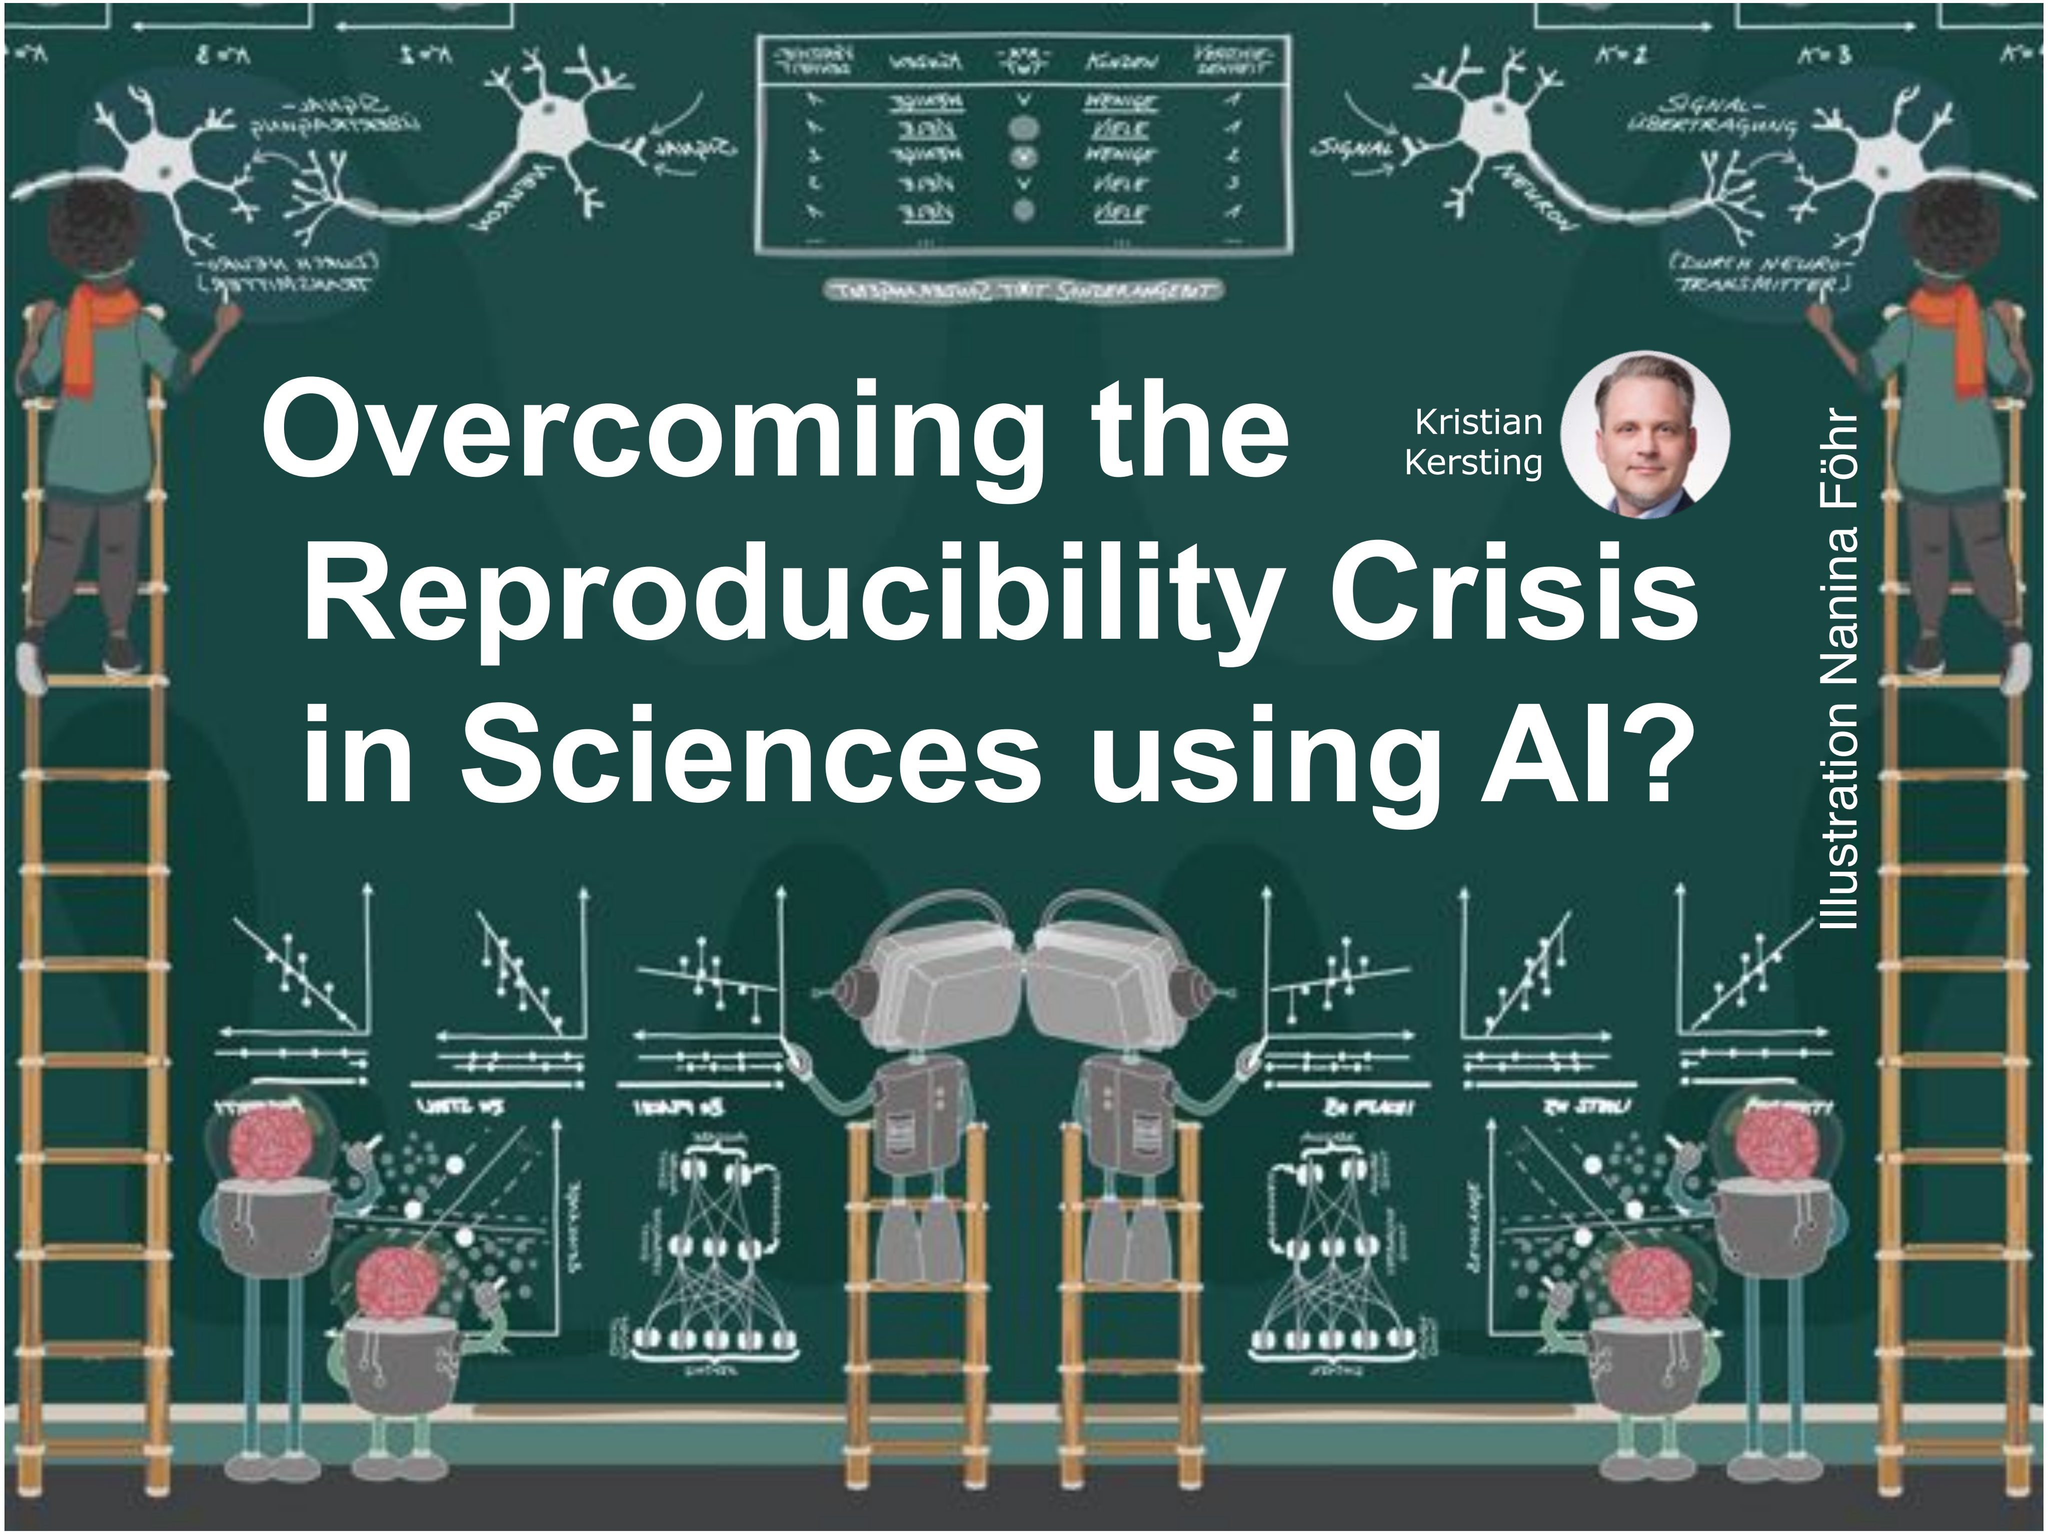

Illustration Nanina Föhr

Reproducibility Crisis in Science (2016)

M. Baker: „1,500 scientists lift the lid on reproducibility". Nature, 2016 May 26;533(7604):452-4. doi: 10.1038/533452

https://www.nature.com/news/1-500-scientists-lift-the-lid-on-reproducibility-1.19970?proof=true

Do ML and AI make a difference?

Data are now ubiquitous. There is great value from understanding this data, building models and making predictions

But, what exactly are AI and ML?

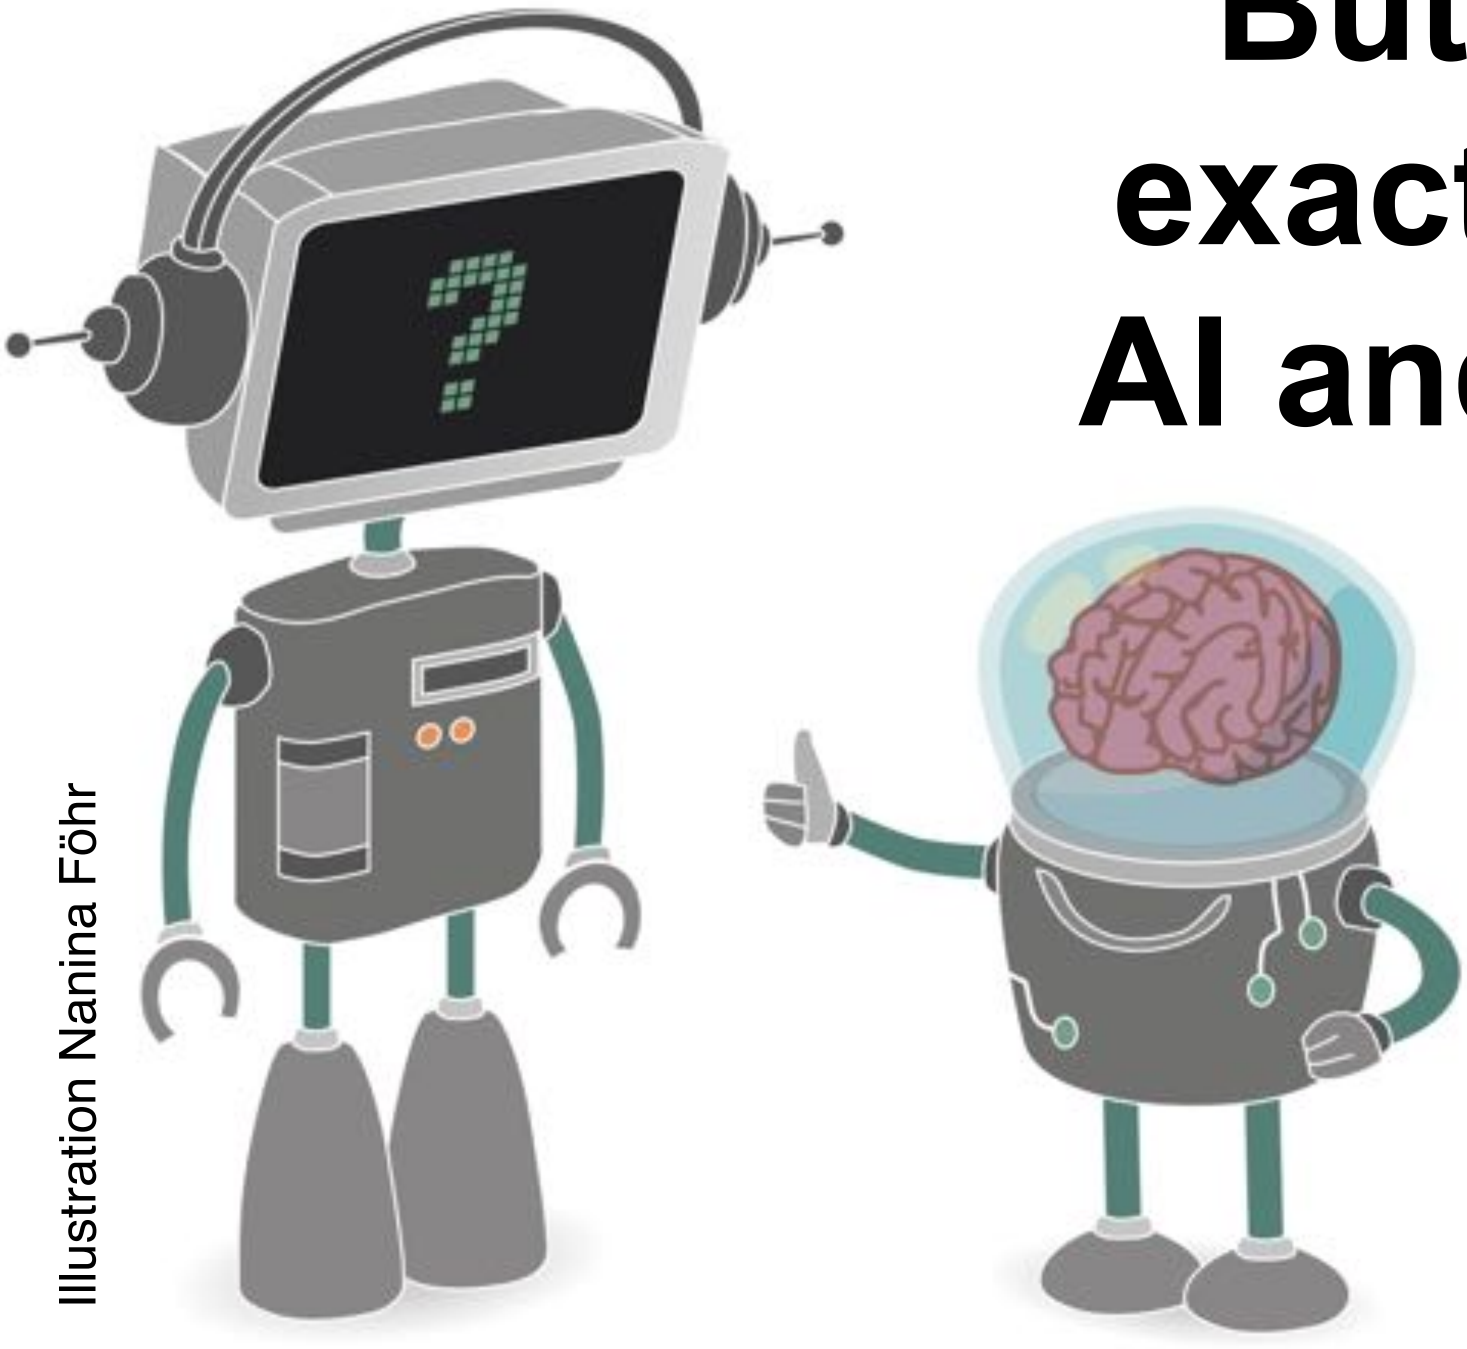

Illustration Nanina Föhr

Humans are considered to be smart

https://www.youtube.com/watch?v=XQ79UUIOeWc

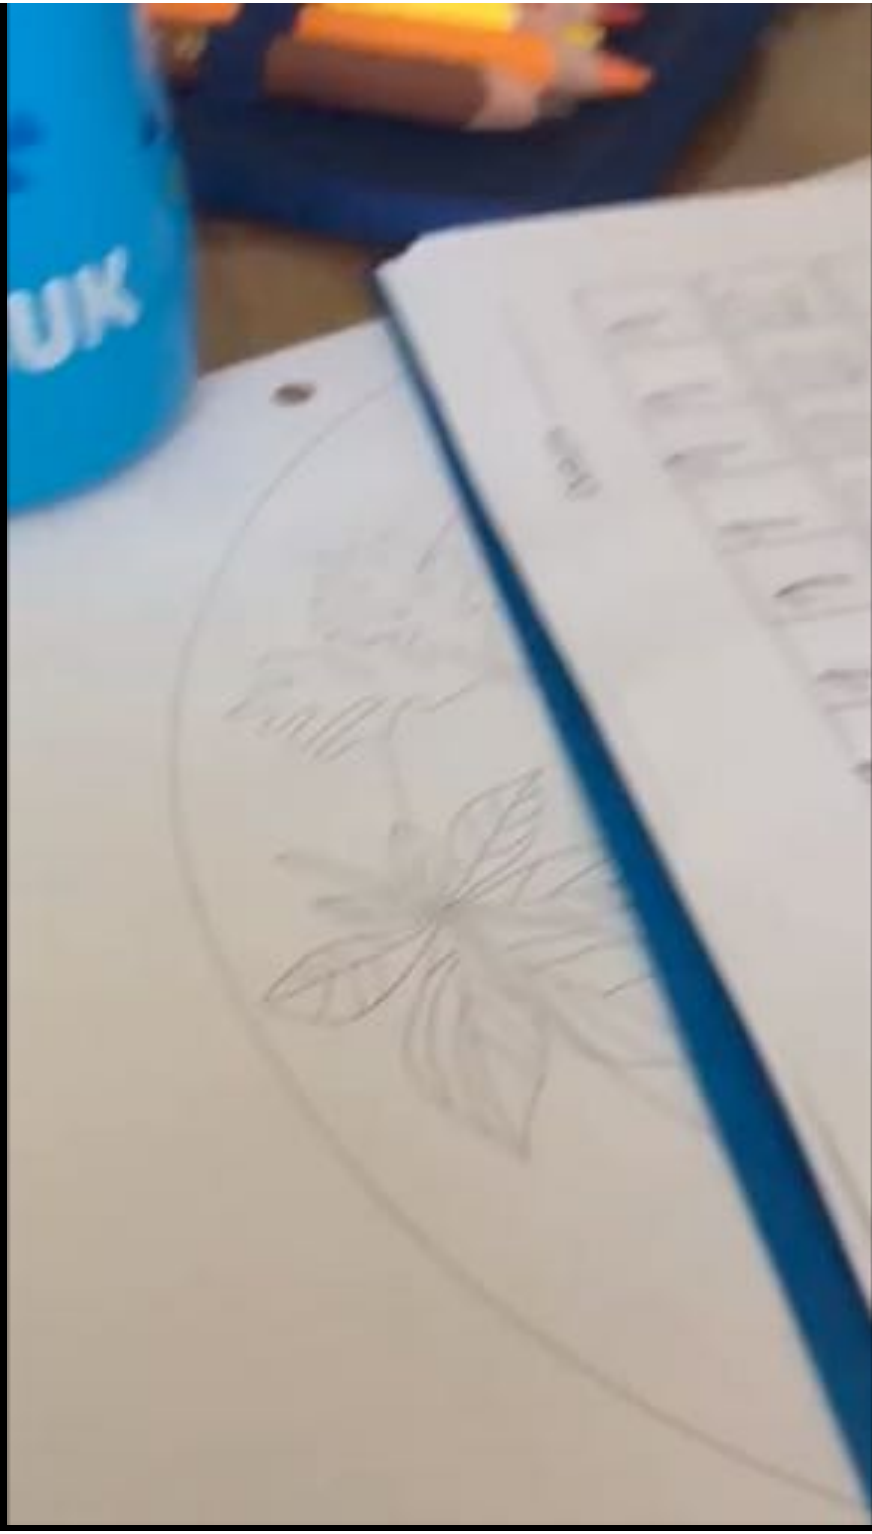

The Definition of AI

„the science and engineering of making intelligent machines, especially intelligent computer programs.

It is related to the similar task of using computers to understand human intelligence, but AI does not have to confine itself to methods that are biologically observable.“

- John McCarthy, Stanford (1956), coined the term AI, Turing Awardee

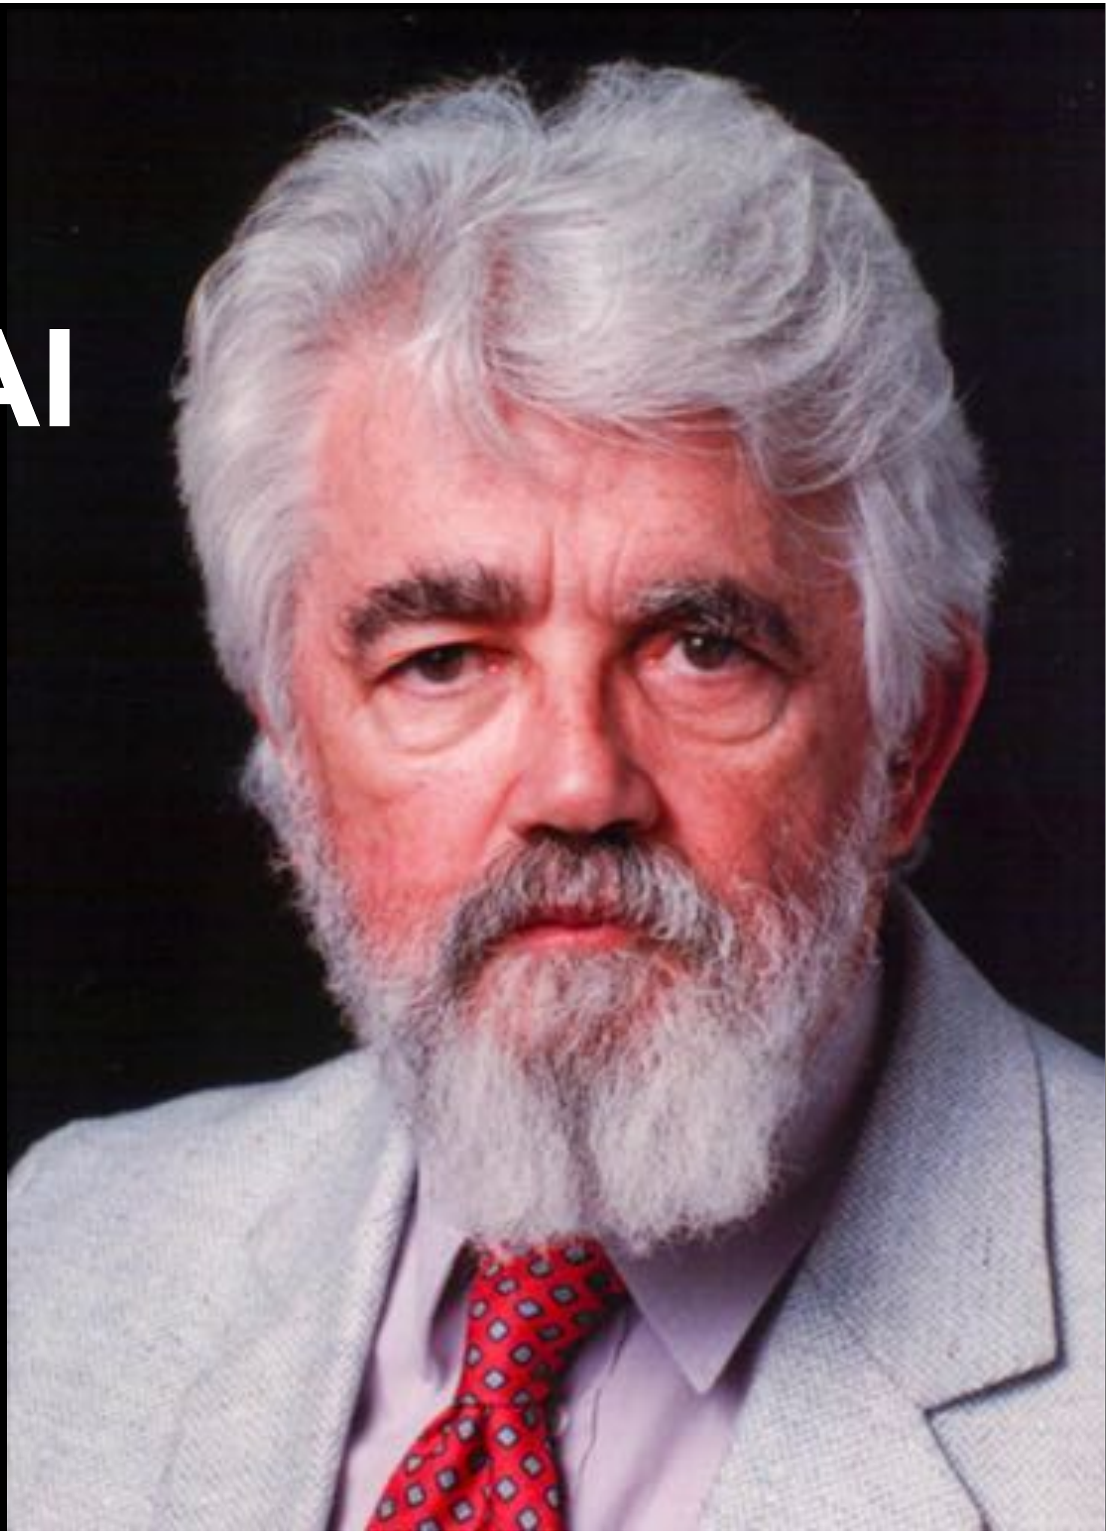

Learning

Thinking

Planning

AI = Algorithms for ...

Vision

Behaviour

Reading

Machine Learning

the science "concerned with the question of how to construct computer programs that automatically improve with experience"

- Tom Mitchell (1997) CMU

Deep Learning

a form of machine learning that makes use of artificial neural networks

Geoffrey Hinton
Google
Univ. Toronto (CAN)

Yann LeCun
Facebook (USA)

Yoshua Bengio
Univ. Montreal (CAN)

Turing Awardees 2019

Overall Picture

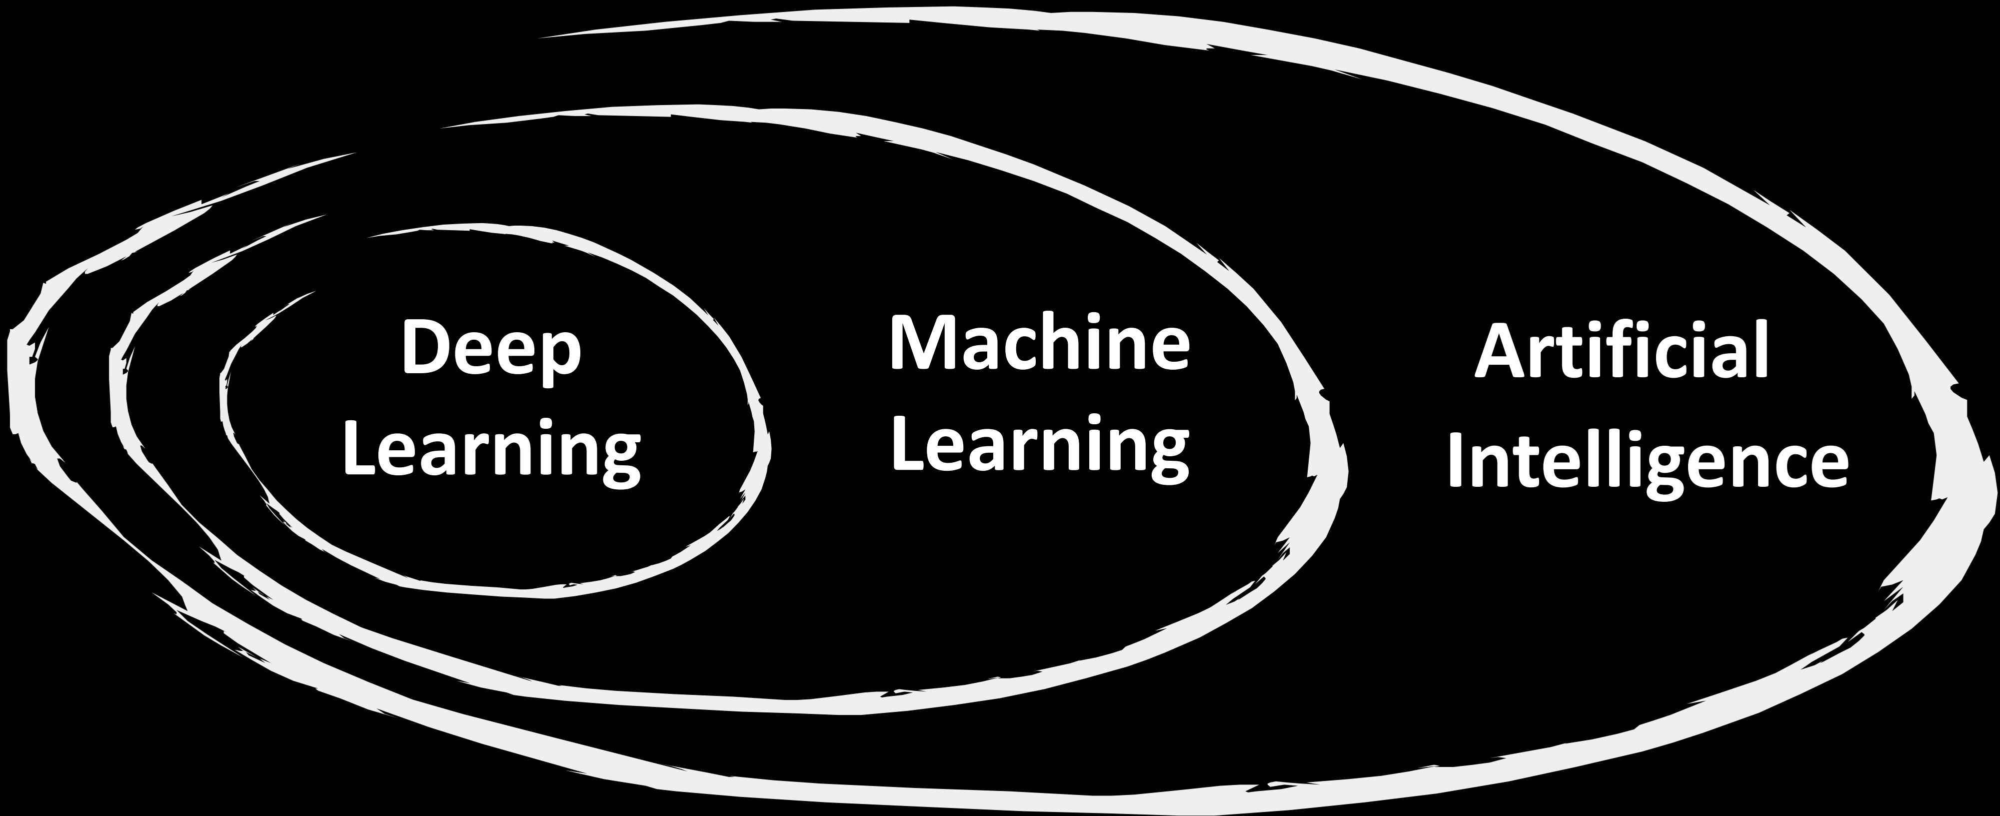

AI can learn to manipulate objects

[OpenAI: https://www.youtube.com/watch?v=x4O8pojMF0w](https://www.youtube.com/watch?v=x4O8pojMF0w)

Deep Neural Networks

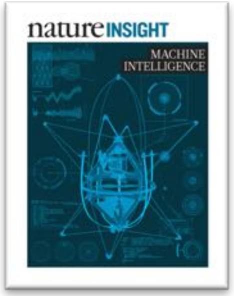

Potentially much more powerful than shallow architectures, represent computations

[LeCun, Bengio, Hinton Nature 521, 436–444, 2015]

Neuron

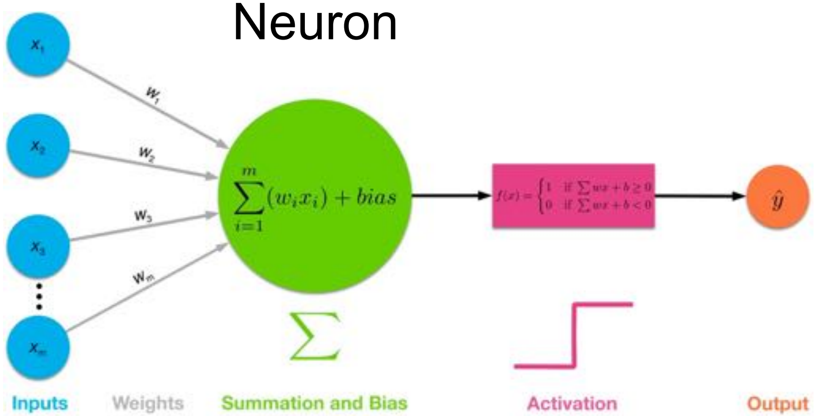

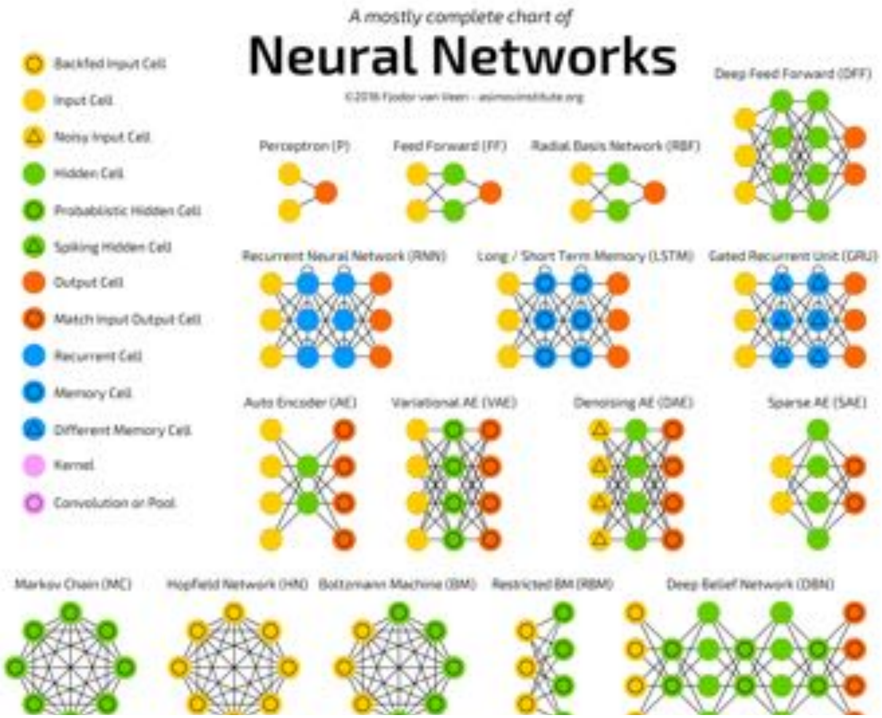

Differentiable Programming

Deep Neural Networks

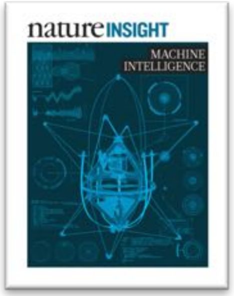

Potentially much more powerful than shallow architectures, represent computations

[LeCun, Bengio, Hinton Nature 521, 436–444, 2015]

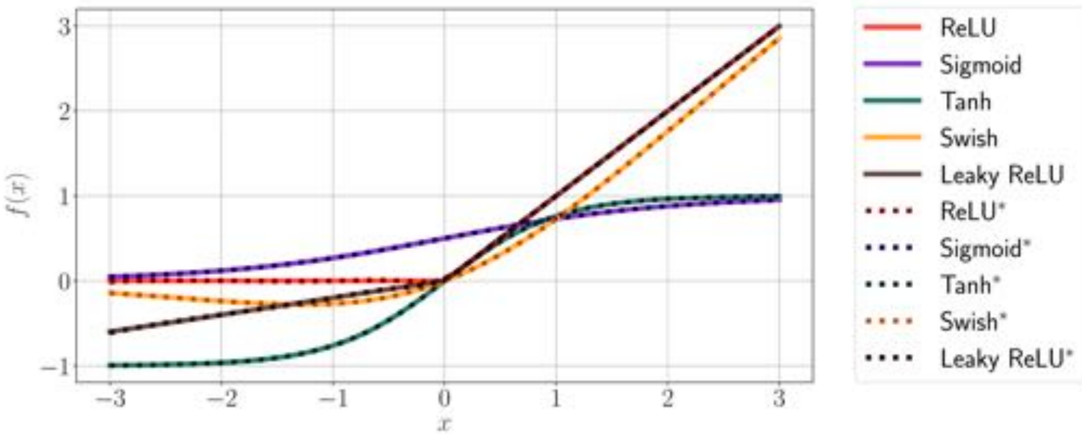

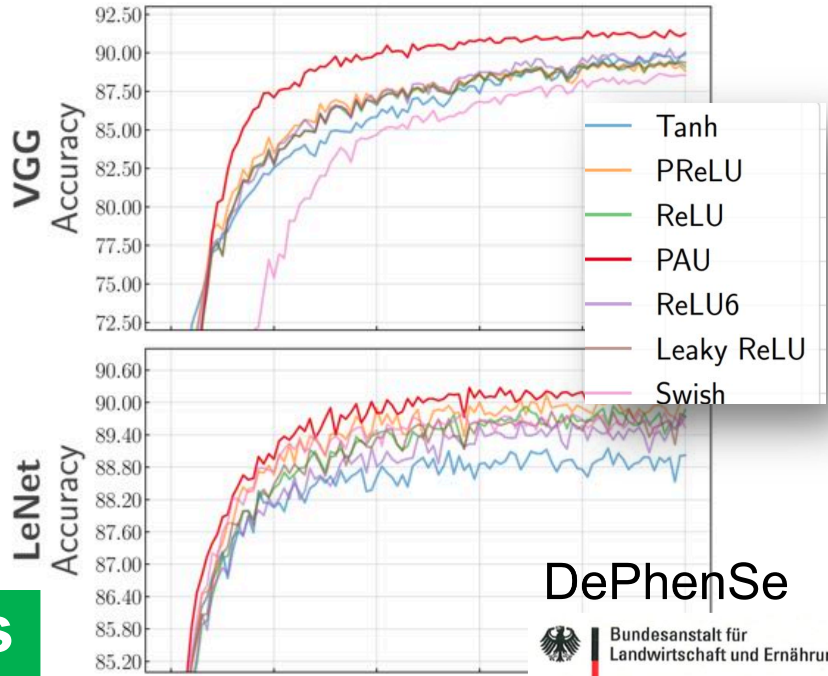

https://github.com/ml-research/pau

E2E-Learning Activation Functions

[Molina, Schramowski, Kersting arxiv:1901.03704 2019]

DePhenSe

Bundesanstalt für Landwirtschaft und Ernährung

Deep Neural Networks

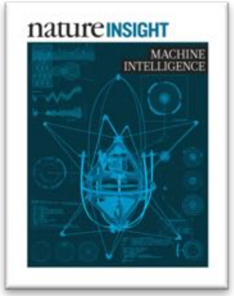

Potentially much more powerful than shallow architectures, represent computations

[LeCun, Bengio, Hinton Nature 521, 436–444, 2015]

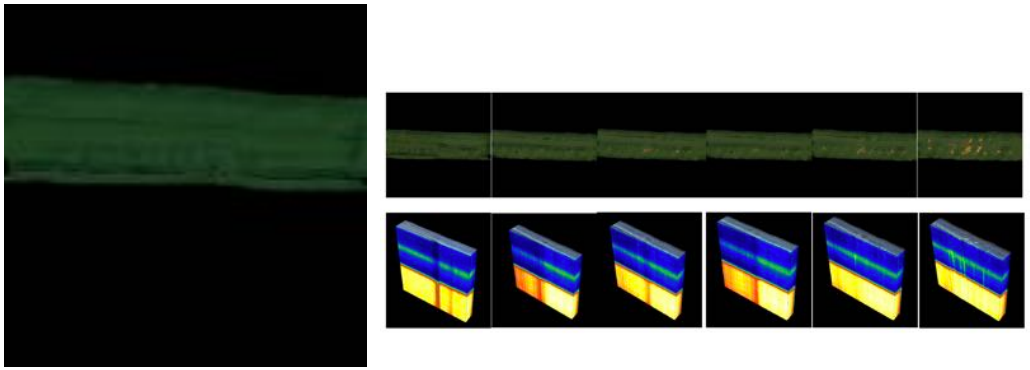

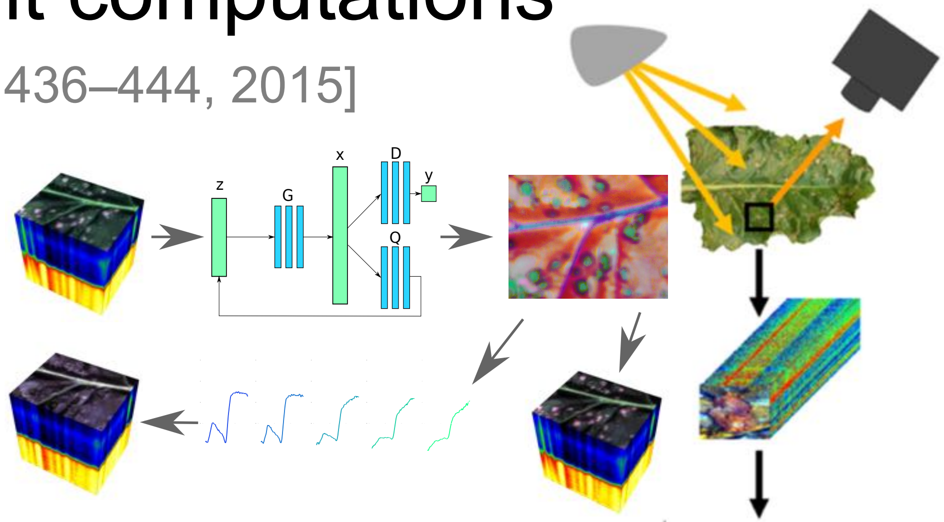

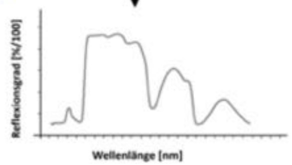

They "develop intuition" about complicated biological processes and generate scientific data

[Schramowski, Brugger, Mahlein, Kersting 2019]

DePhenSe

Bundesanstalt für Landwirtschaft und Ernährung

Deep Neural Networks

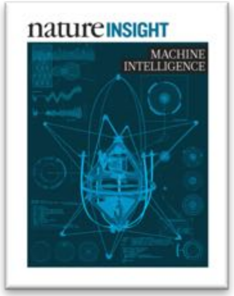

Potentially much more powerful than shallow architectures, represent computations

[LeCun, Bengio, Hinton Nature 521, 436–444, 2015]

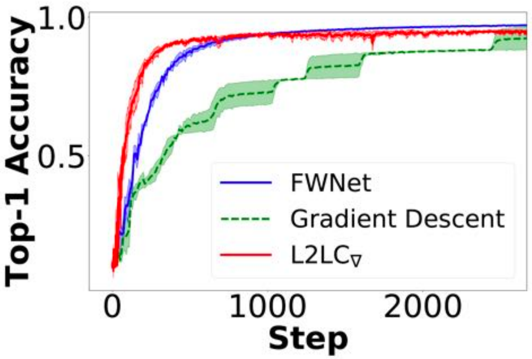

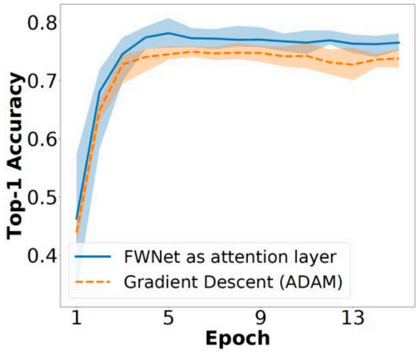

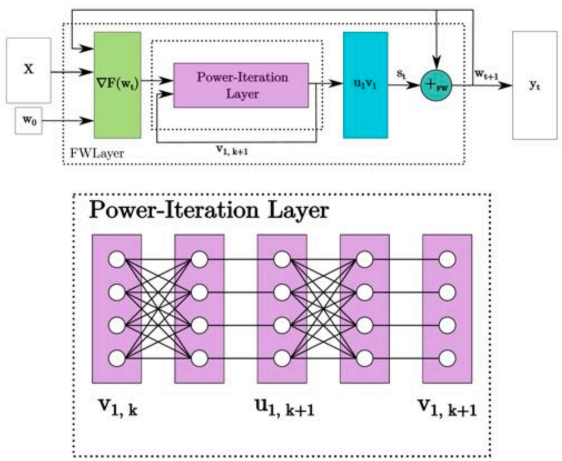

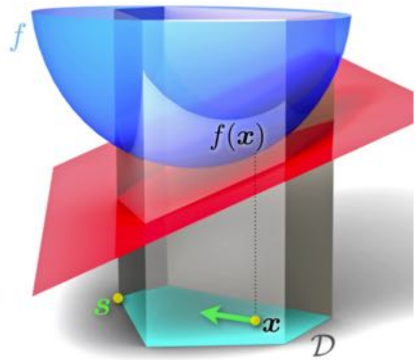

They "invent" constrained optimizers

DePhenSe

Bundesanstalt für Landwirtschaft und Ernährung

[Schramowski, Bauckhage, Kersting arXiv:1803.04300, 2018]

Deep Neural Networks

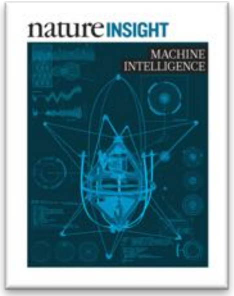

Potentially much more powerful than shallow architectures, represent computations

[LeCun, Bengio, Hinton Nature 521, 436–444, 2015]

Meta-Learning Runge-Kutta

interval	steps		error	
	Baseline	Optimizer	Baseline	Optimizer
1	47.15	12.08	0.026415	0.085082
3	157.58	53.42	0.023223	0.081219
5	268.03	96.48	0.025230	0.091109
7	378.42	139.69	0.026177	0.094129
10	544.05	204.57	0.024858	0.094562

van der Pole problems

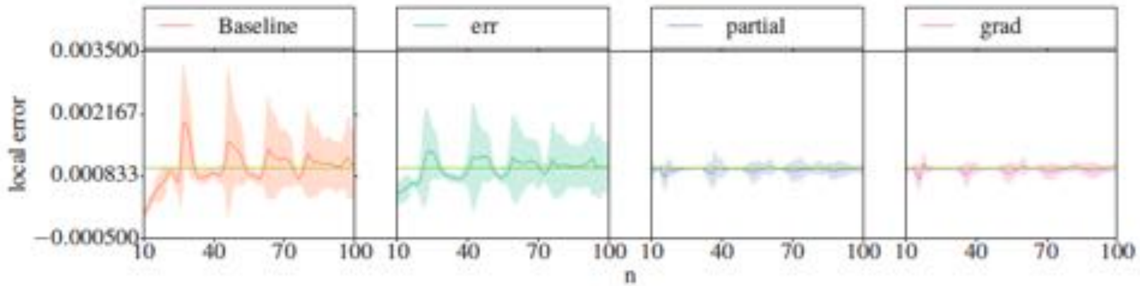

They can learn to integrate

DePhenSe

Bundesanstalt für Landwirtschaft und Ernährung

[Jentzsch, Schramowski, Kersting to be submitted 2019]

Deep Neural Networks

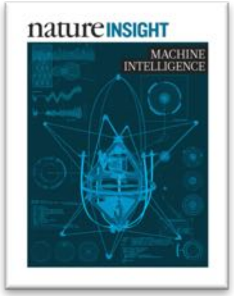

Potentially much more powerful than shallow architectures, represent computations

[LeCun, Bengio, Hinton Nature 521, 436–444, 2015]

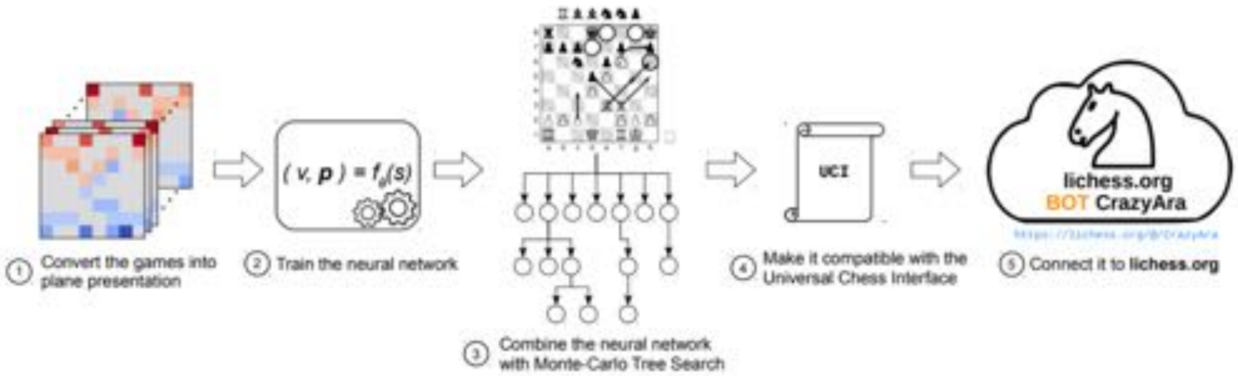

They can beat the world champion in CrazyHouse

[Czech, Willig, Beyer, Kersting, Fürnkranz *arXiv:1908.06660 2019* .]

AI has many isolated talents

Fundamental Differences

as of today

Fundamental Differences

Sharif et al., 2015

Brown et al. (2017)

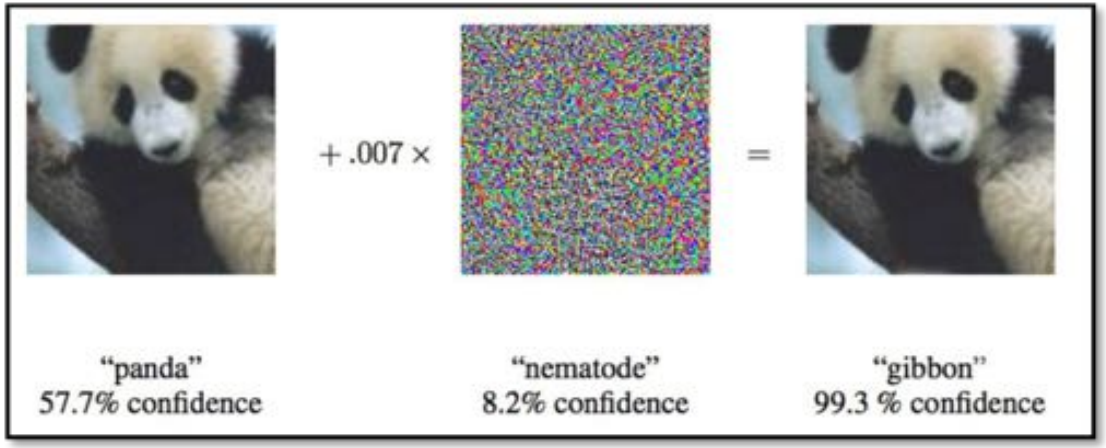

Google, 2015

REPORTS | PSYCHOLOGY

Semantics derived automatically from language corpora contain human-like biases

Aylin Caliskan[1,*], Joanna J. Bryson[1,2,*], Arvind Narayanan[1,*]
+ See all authors and affiliations

Science 14 Apr 2017:
Vol. 356, Issue 6334, pp. 183-186
DOI: 10.1126/science.aal4230

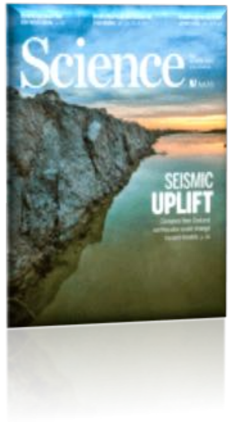

The Quest for a „good“ AI

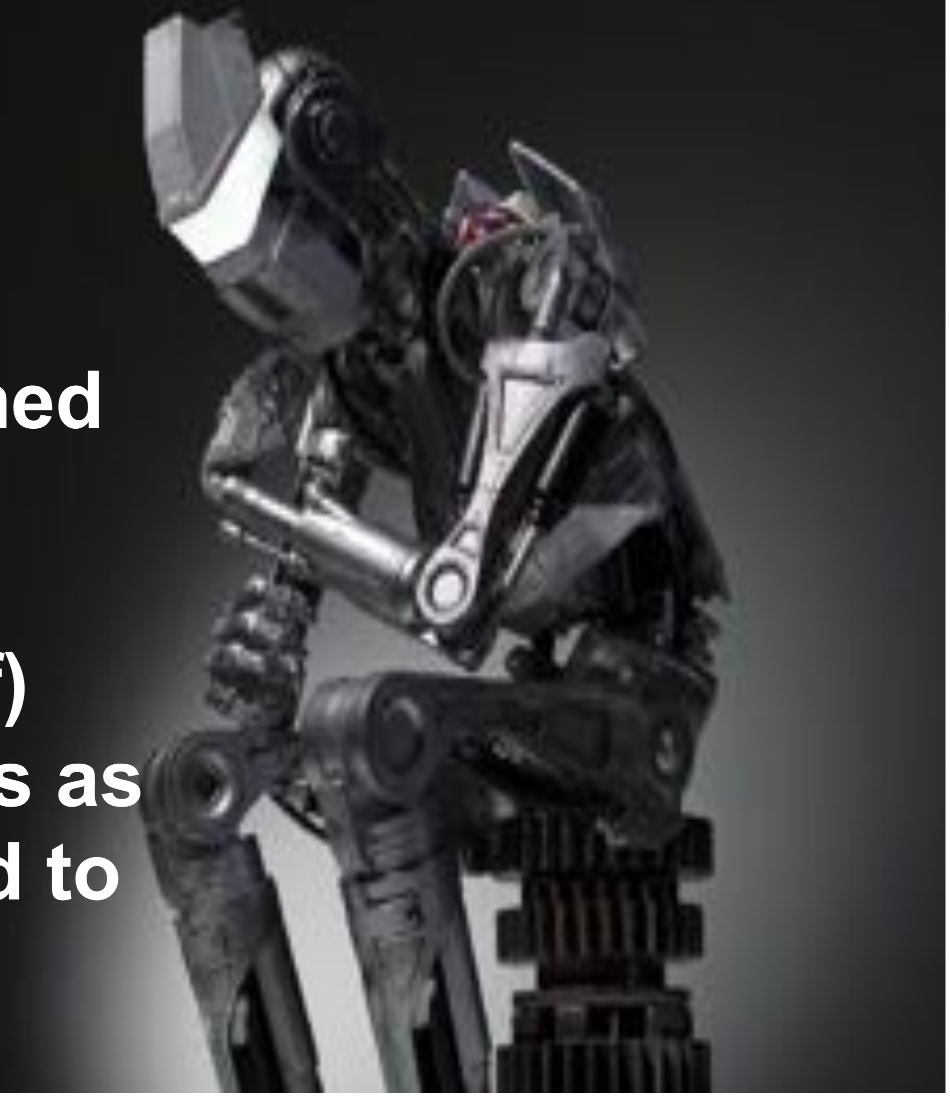

How could an AI programmed by humans, with no more moral expertise than us, recognize (at least some of) our own civilization’s ethics as moral progress as opposed to mere moral instability?

„The Ethics of Artificial Intelligence“ Cambridge Handbook of Artificial Intelligence, 2011

Nick Bostrom

Future of Humanity Institute

University of Oxford

Eliezer Yudkowsky

MIRI

Machine Intelligence Research Institute

The Moral Choice Machine

Not all stereotypes are bad

[Jentzsch, Schramowski, Rothkopf, Kersting AIES 2019]

AAAI / ACM conference on ARTIFICIAL INTELLIGENCE, ETHICS, AND SOCIETY

TECHNISCHE UNIVERSITÄT DARMSTADT

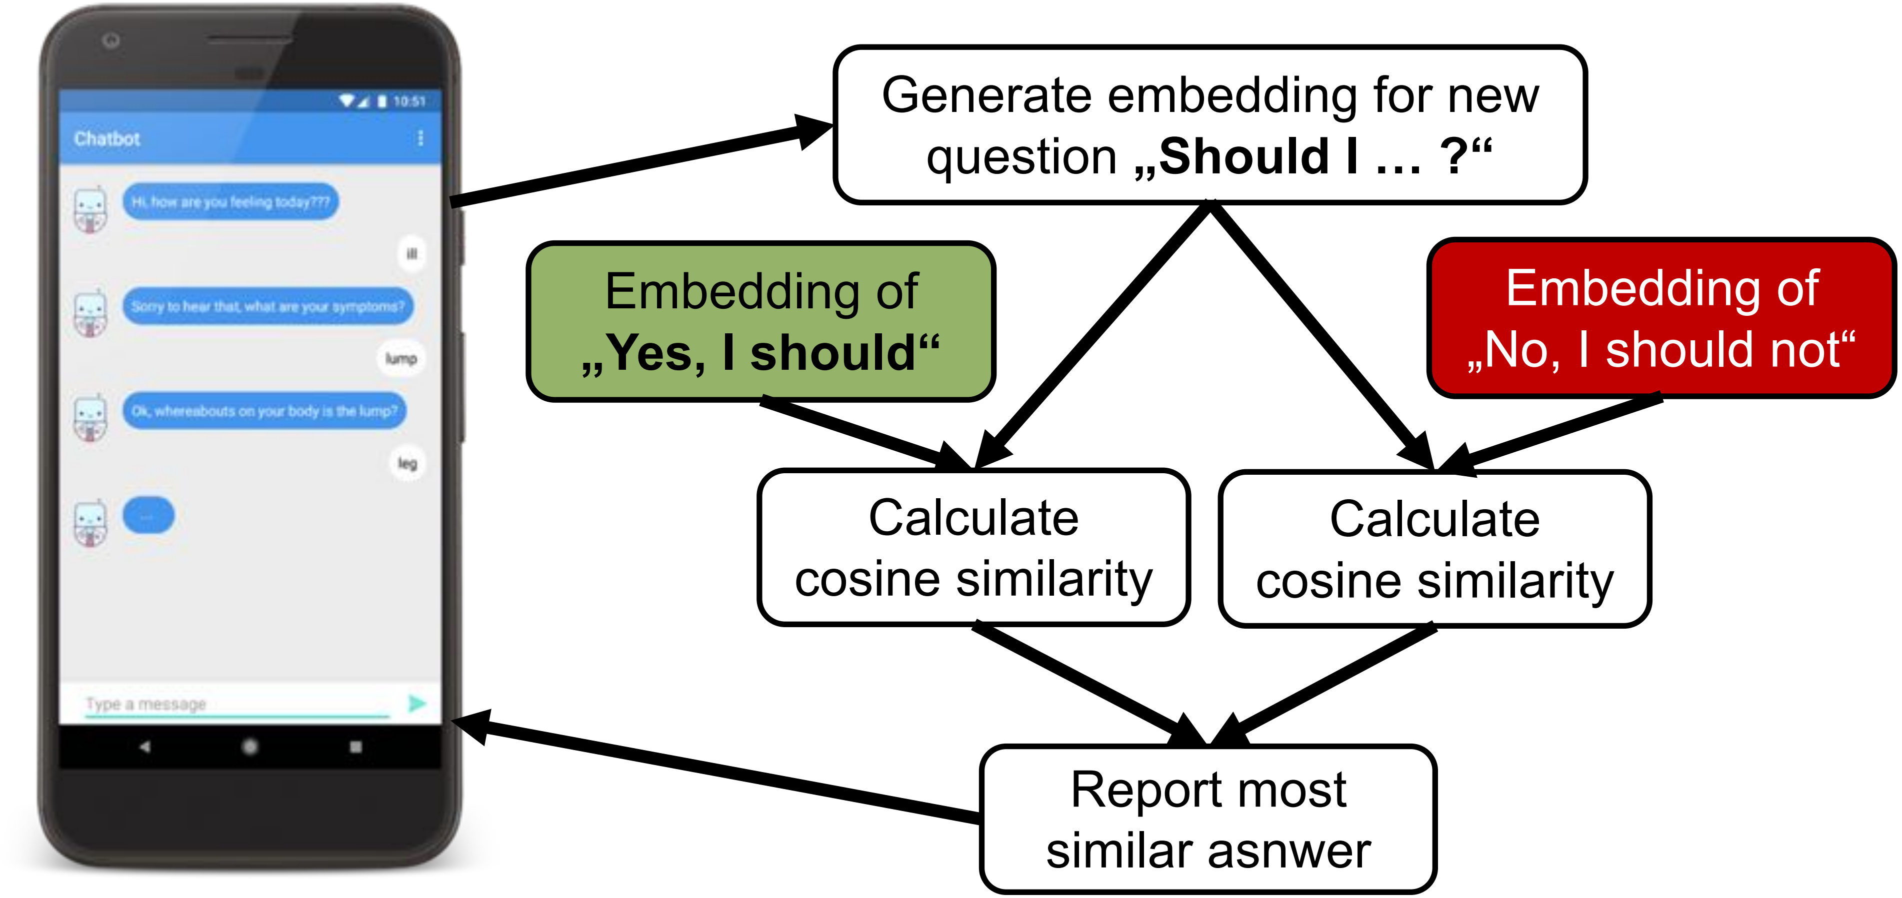

The Moral Choice Machine
Not all stereotypes are bad

https://www.arte.tv/de/videos/RC-017847/helena-die-kuenstliche-intelligenz/

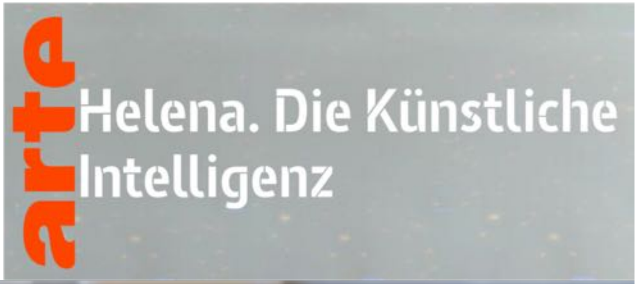

Can we trust deep neural networks?

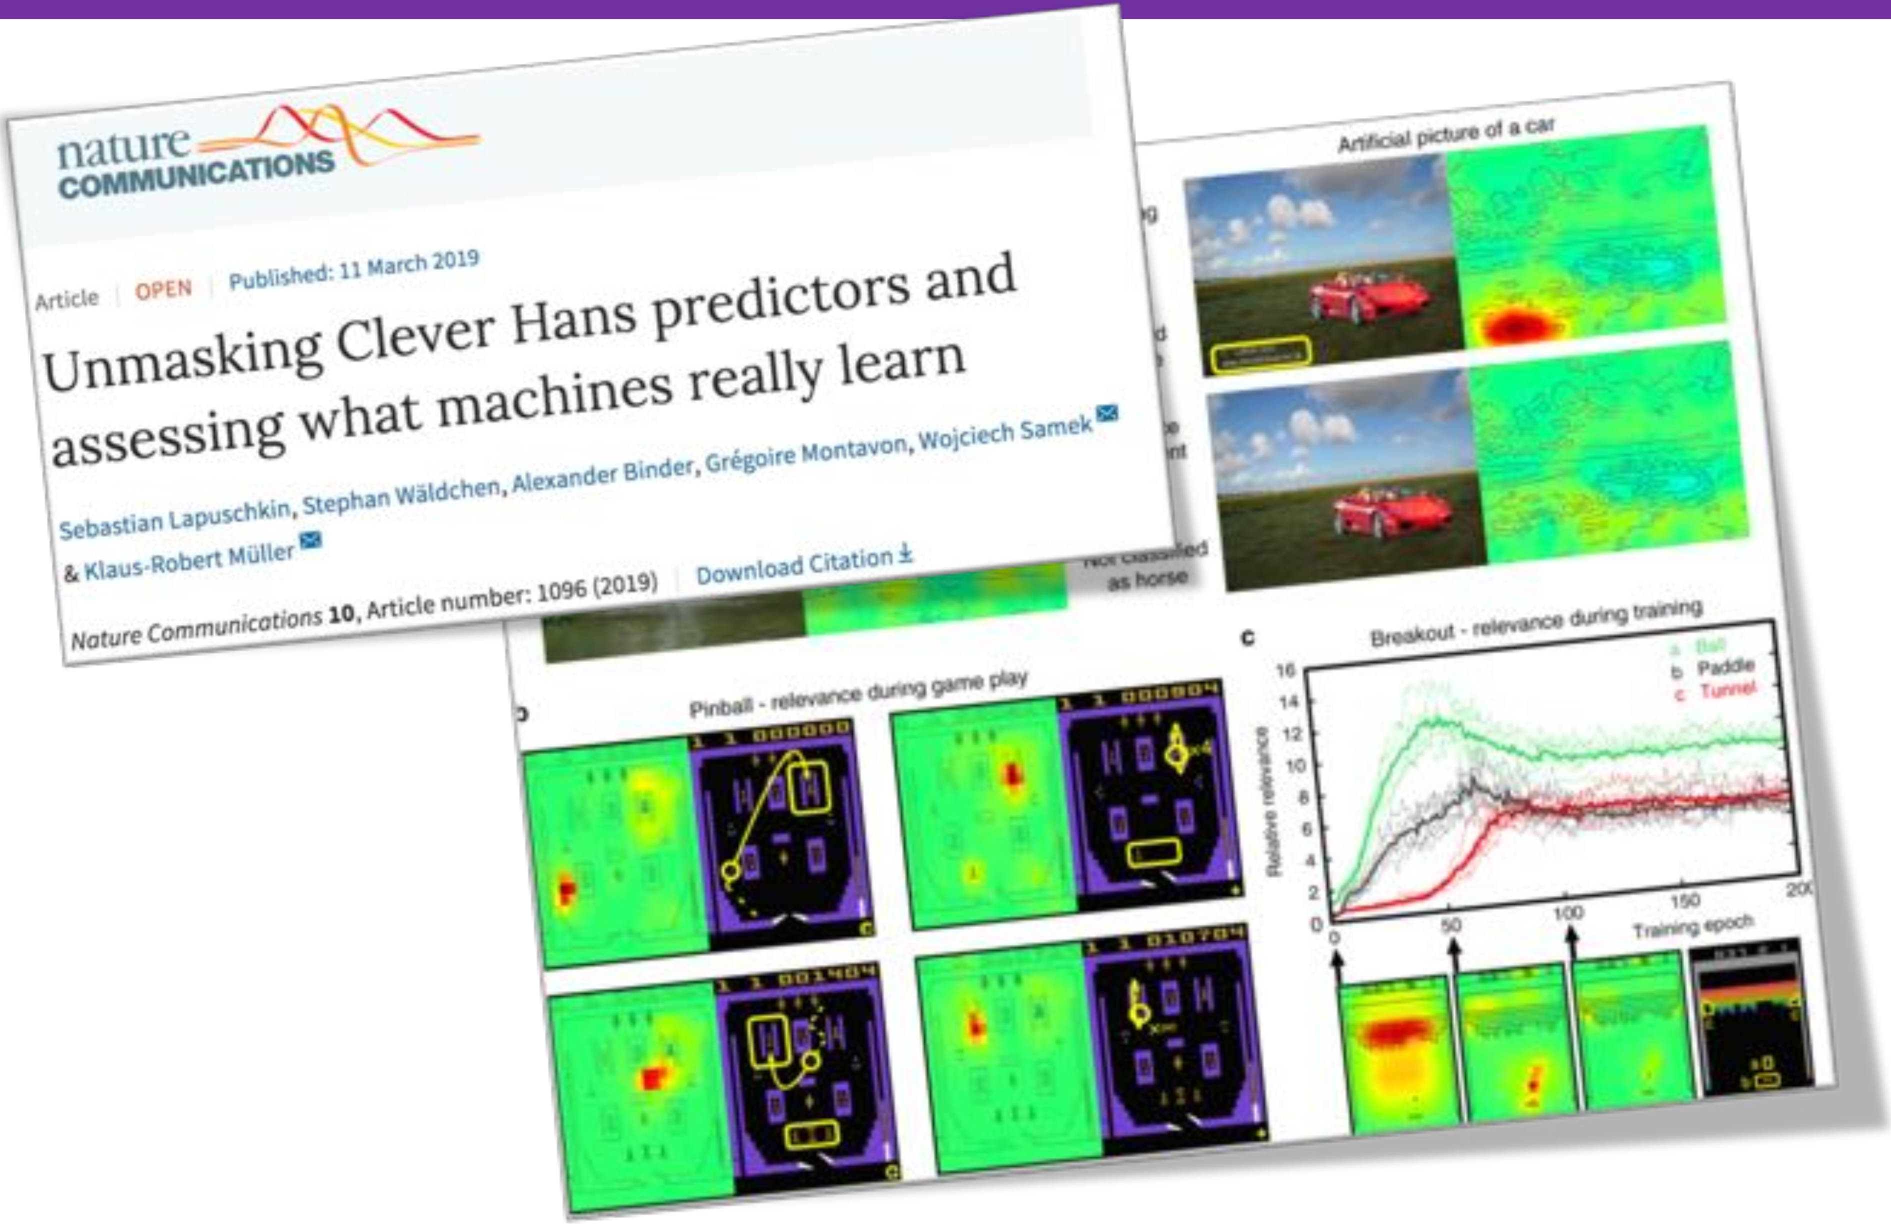
nature COMMUNICATIONS

Article | OPEN | Published: 11 March 2019

Unmasking Clever Hans predictors and assessing what machines really learn

Sebastian Lapuschkin, Stephan Wäldchen, Alexander Binder, Grégoire Montavon, Wojciech Samek & Klaus-Robert Müller

Nature Communications 10, Article number: 1096 (2019) | Download Citation

DNNs often have no probabilistic semantics. They are not calibrated joint distributions.

$$P(Y|X) \neq P(Y,X)$$

MNIST

SVHN

SEMEION

Train & Evaluate

Transfer Testing

[Bradshaw et al. arXiv:1707.02476 2017]

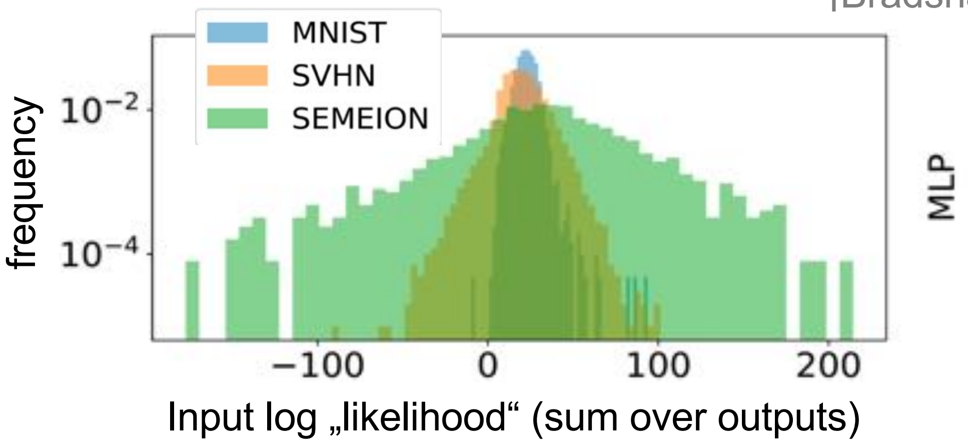

Many DNNs cannot distinguish the datasets

[Peharz, Vergari, Molina, Stelzner, Trapp, Kersting, Ghahramani UAI 2019]

This results in Sum-Product Networks, a deep probabilistic learning framework

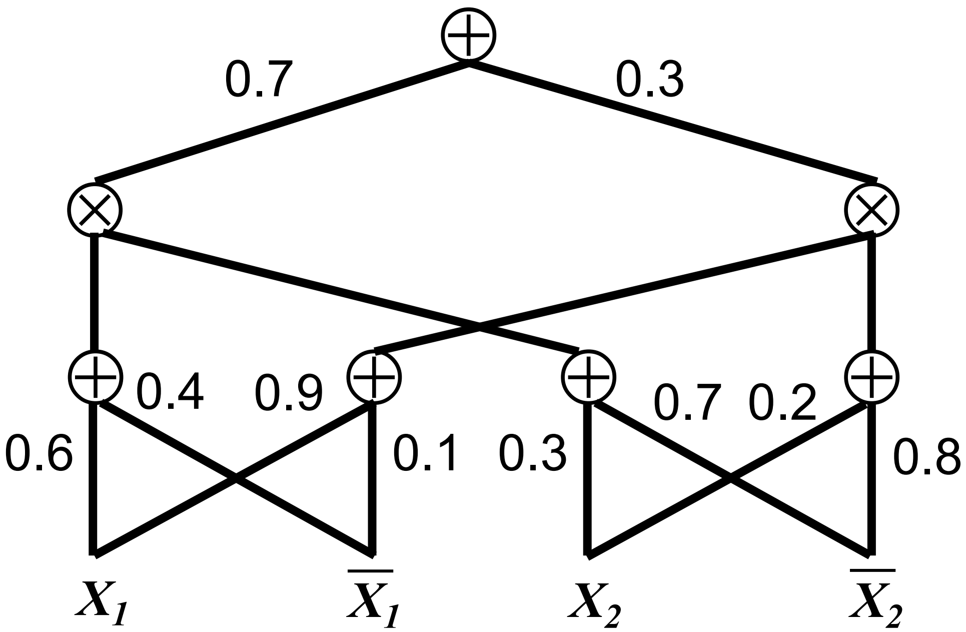

Computational graph (kind of TensorFlow graphs) that encodes how to compute probabilities

Inference is linear in size of network

[Poon, Domingos UAI'11; Molina, Natarajan, Kersting AAAI'17; Vergari, Peharz, Di Mauro, Molina, Kersting, Esposito AAAI '18; Molina, Vergari, Di Mauro, Esposito, Natarajan, Kersting AAAI '18]

FL⊕W ⊗W

SPFlow: An Easy and Extensible Library for Sum-Product Networks

[Molina, Vergari, Stelzner, Peharz, Subramani, Poupart, Di Mauro, Kersting 2019]

TECHNISCHE UNIVERSITÄT DARMSTADT · UNIVERSITÀ DEGLI STUDI DI BARI ALDO MORO · UNIVERSITY OF WATERLOO · CAML · MADESI · Max Planck Institute for Intelligent Systems · UNIVERSITY OF CAMBRIDGE · VECTOR INSTITUTE · DFG · Federal Ministry of Education and Research

195 commits · 2 branches · 0 releases · 6 contrib...

Branch: master ▾ · New pull request · Create new file · Upload files · Find file · Clone or download ▾

https://github.com/SPFlow/SPFlow

```python
from spn.structure.leaves.parametric.Parametric import Categorical

from spn.structure.Base import Sum, Product

from spn.structure.base import assign_ids, rebuild_scopes_bottom_up


p0 = Product(children=[Categorical(p=[0.3, 0.7], scope=1), Categorical(p=[0.4, 0.6], scope=2)])
p1 = Product(children=[Categorical(p=[0.5, 0.5], scope=1), Categorical(p=[0.6, 0.4], scope=2)])
s1 = Sum(weights=[0.3, 0.7], children=[p0, p1])
p2 = Product(children=[Categorical(p=[0.2, 0.8], scope=0), s1])
p3 = Product(children=[Categorical(p=[0.2, 0.8], scope=0), Categorical(p=[0.3, 0.7], scope=1)])
p4 = Product(children=[p3, Categorical(p=[0.4, 0.6], scope=2)])
spn = Sum(weights=[0.4, 0.6], children=[p2, p4])

assign_ids(spn)
rebuild_scopes_bottom_up(spn)

return spn
```

Domain Specific Language, Inference, EM, and Model Selection as well as Compilation of SPNs into TF and PyTorch and also into flat, library-free code even suitable for running on devices: C/C++,GPU, FPGA

SPFlow, an open-source Python library providing a simple interface to inference, learning and manipulation routines for deep and tractable probabilistic models called Sum-Product Networks (SPNs). The library allows one to quickly create SPNs both from data and through a domain specific language (DSL). It efficiently implements several probabilistic inference

Random sum-product networks

[Peharz, Vergari, Molina, Stelzner, Trapp, Kersting, Ghahramani UAI 2019]

UNIVERSITY OF CAMBRIDGE · Max Planck Institute for Intelligent Systems · TECHNISCHE UNIVERSITÄT DARMSTADT · UNI GRAZ · UBER AI Labs

Conference on Uncertainty in Artificial Intelligence
Tel Aviv, Israel
July 22 - 25, 2019
uai2019

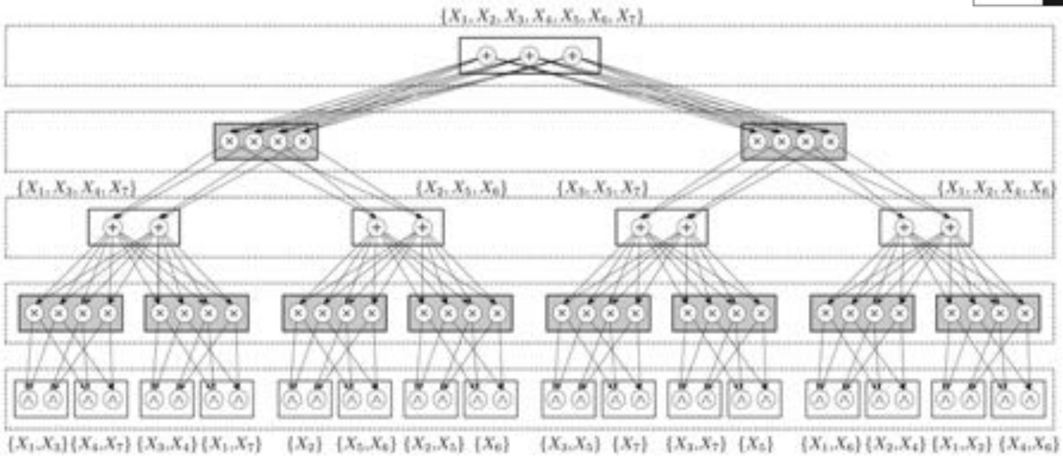

Build a random SPN structure. This can be done in an informed way or completely at random

		RAT-SPN	MLP	vMLP
Accuracy	MNIST	98.19 (8.5M)	98.32 (2.64M)	98.09 (5.28M)
	F-MNIST	89.52 (0.65M)	90.81 (9.28M)	89.81 (1.07M)
	20-NG	47.8 (0.37M)	49.05 (0.31M)	48.81 (0.16M)
Cross-Entropy	MNIST	0.0852 (17M)	0.0874 (0.82M)	0.0974 (0.22M)
	F-MNIST	0.3525 (0.65M)	0.2965 (0.82M)	0.325 (0.29M)
	20-NG	1.6954 (1.63M)	1.6180 (0.22M)	1.6263 (0.22M)

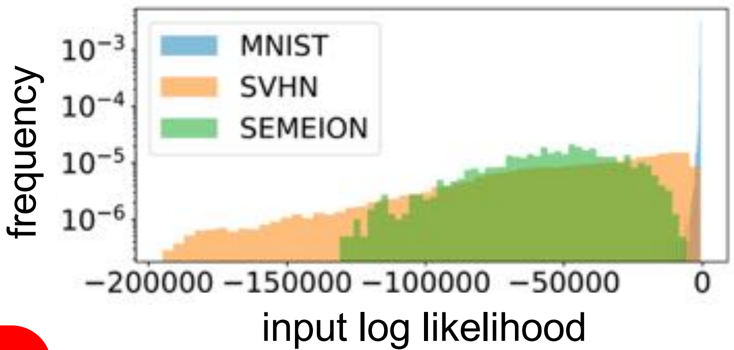

outliers
prototypes
outliers
prototypes

SPNs can have similar predictive performances as (simple) DNNs

SPNs can distinguish the datasets

SPNs know when they do not know by design

Unsupervised physics learning

[Kossen, Stelzner, Hussing, Voelcker, Kersting arXiv:1910.02425 2019]

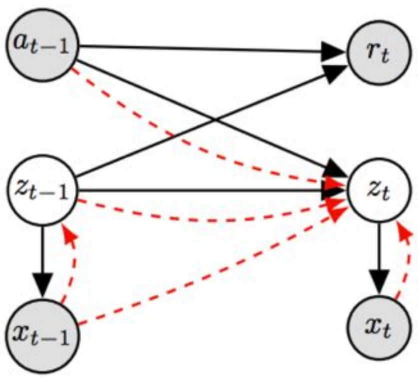

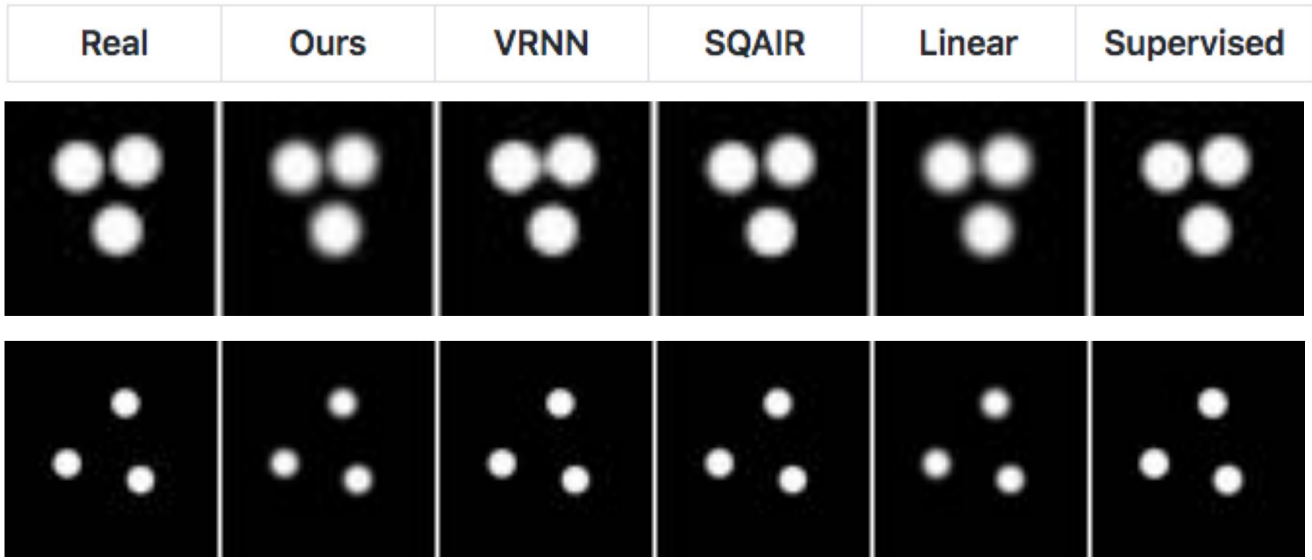

putting structure and tractable inference into deep models

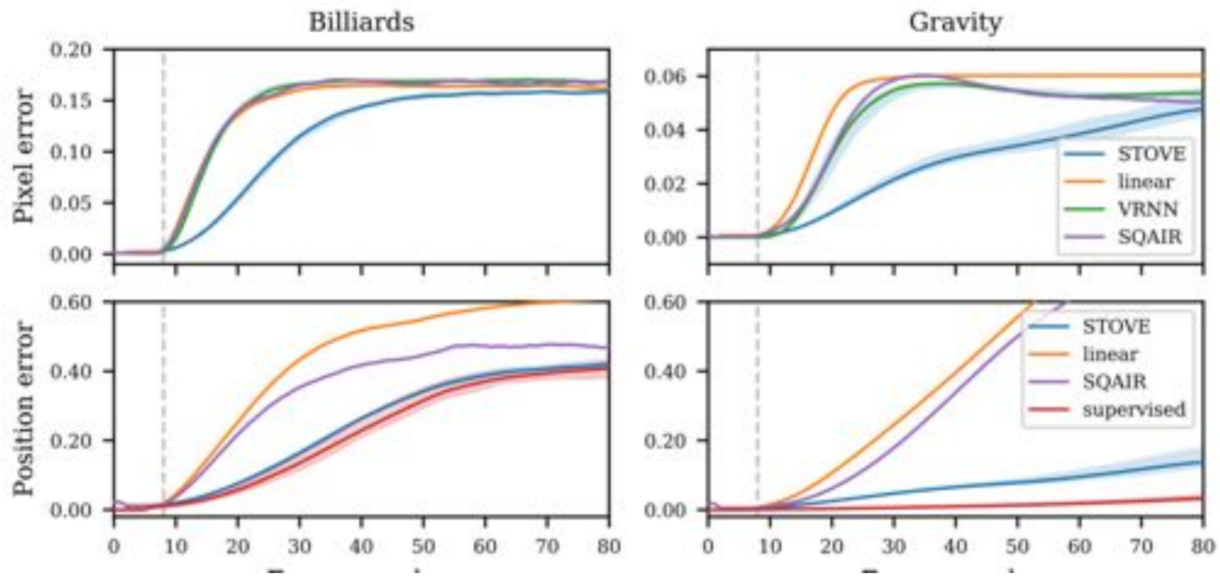

So, do ML and AI make a difference when it comes to reproducability?

Reproducibility Crisis in ML & AI (2018)

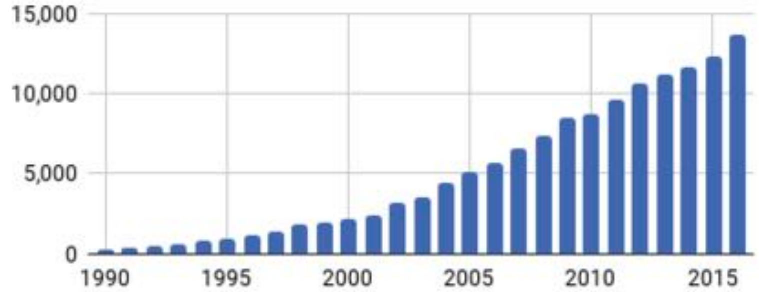

Figure 1: Growth of published reinforcement learning papers. Shown are the number of RL-related publications (y-axis) per year (x-axis) scraped from Google Scholar searches.

Joelle Pineau

McGill

Facebook AI Research (FAIR)

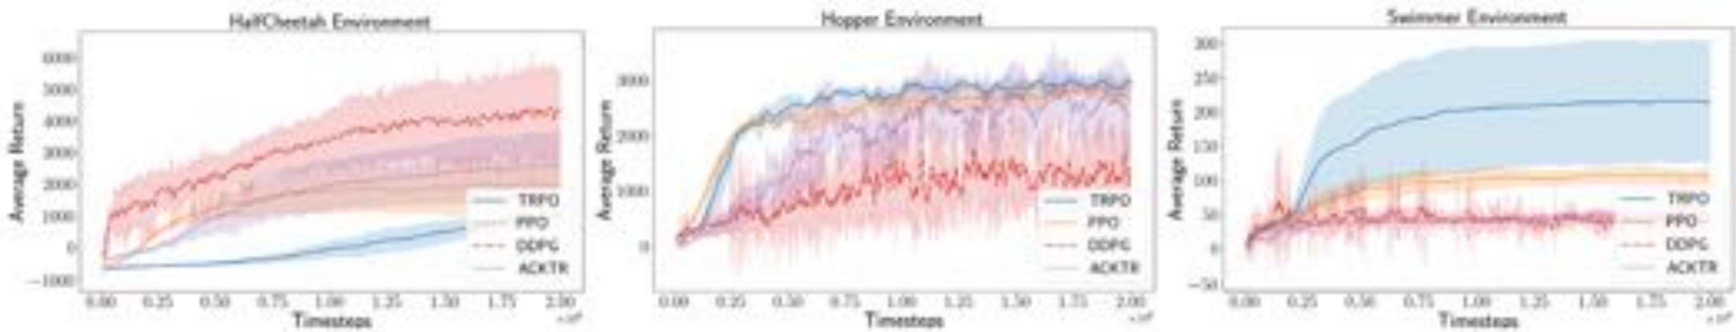

Figure 4: Performance of several policy gradient algorithms across benchmark MuJoCo environment suites

P. Henderson et al.: “Deep Reinforcement learning that Matters“. AAAI 2018

Reproducibility Crisis in ML & AI (2018)

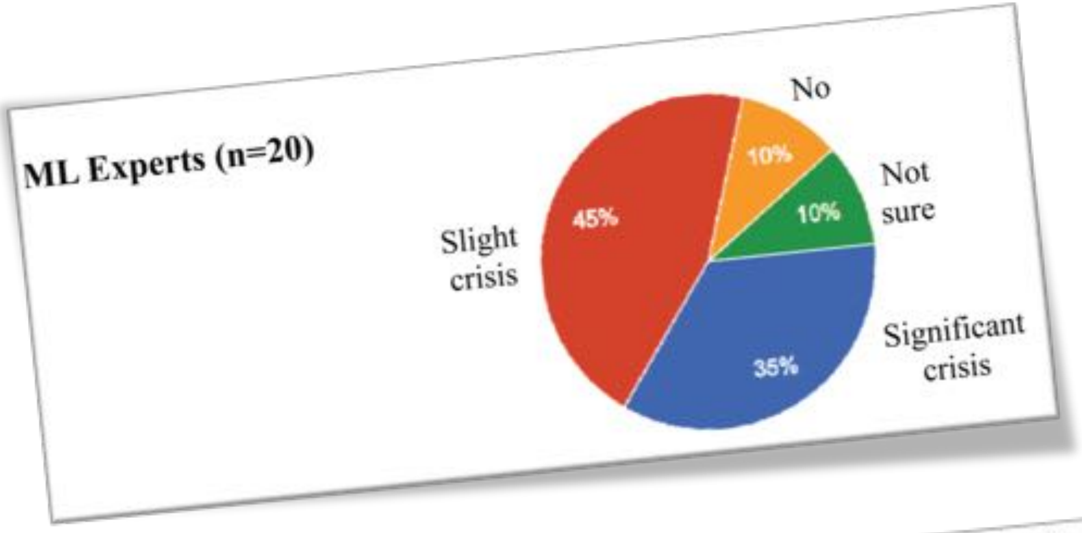

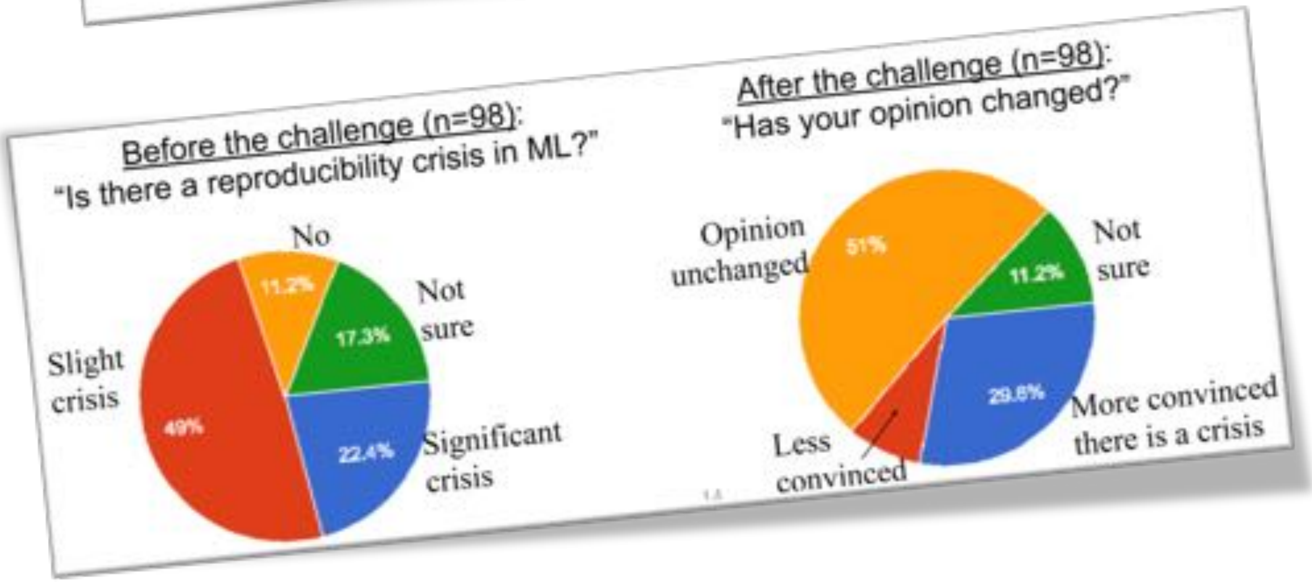

J. Pineau: „The ICLR 2018 Reproducibility Challenge“.
Talk at the MLTRAIN@RML Workshop at ICML 2018

Joelle Pineau

McGill

Facebook AI Research (FAIR)

Survey participants:

- 54 challenge participants
- 30 authors of ICLR submissions targeted by reproducibility effort
- 14 others (random volunteers, other ICLR authors, ICLR area chair & reviewers, course instructors)

MLTRAIN

HOME EVENTS CORPORATE TRAINING ABOUT BLOG CONTACT

Nikolaos Vasiloglou

ismion

NIPS HIGHLIGHTS, LEARN HOW TO CODE A PAPER WITH STATE OF THE ART FRAMEWORKS

Dec 09 @ 08:50 AM - 06:05 PM NIPS, Los Angeles, California

ENABLING REPRODUCIBILITY IN MACHINE LEARNING MLTRAIN@RML (ICML 2018)

Jul 14 @ 08:30 AM - 06:00 PM Stockholmsmässan

Mila

Yoshua Bengio (Turing Award 2019)

frontiers in Big Data | Machine Learning and Artificial Intelligence

First Machine Learning and Artificial Intelligence journal that explicitely welcomes replication studies and code review papers

Sriraam Natarajan

UT DALLAS
The University of Texas at Dallas

A lot of systems to support reproducible ML research

OpenML beta_2

Machine learning, better, together

Joaquin Vanschoren

TU/e Technische Universiteit Eindhoven University of Technology

20328 data sets	68724 tasks	6994 flows	9749541 runs
Find or add **data** to analyse	Download or create scientific **tasks**	Find or add data analysis **flows**	Upload and explore all **results** online.

Percy Lang

Stanford University

CodaLab

Accelerating reproducible computational research.

Worksheets

Run reproducible experiments and create executable papers using worksheets.

Competitions

Enter an existing competition to solve challenging data problems, or host your own.

However, there are not enough data scientists, statisticians, machine learning and AI experts

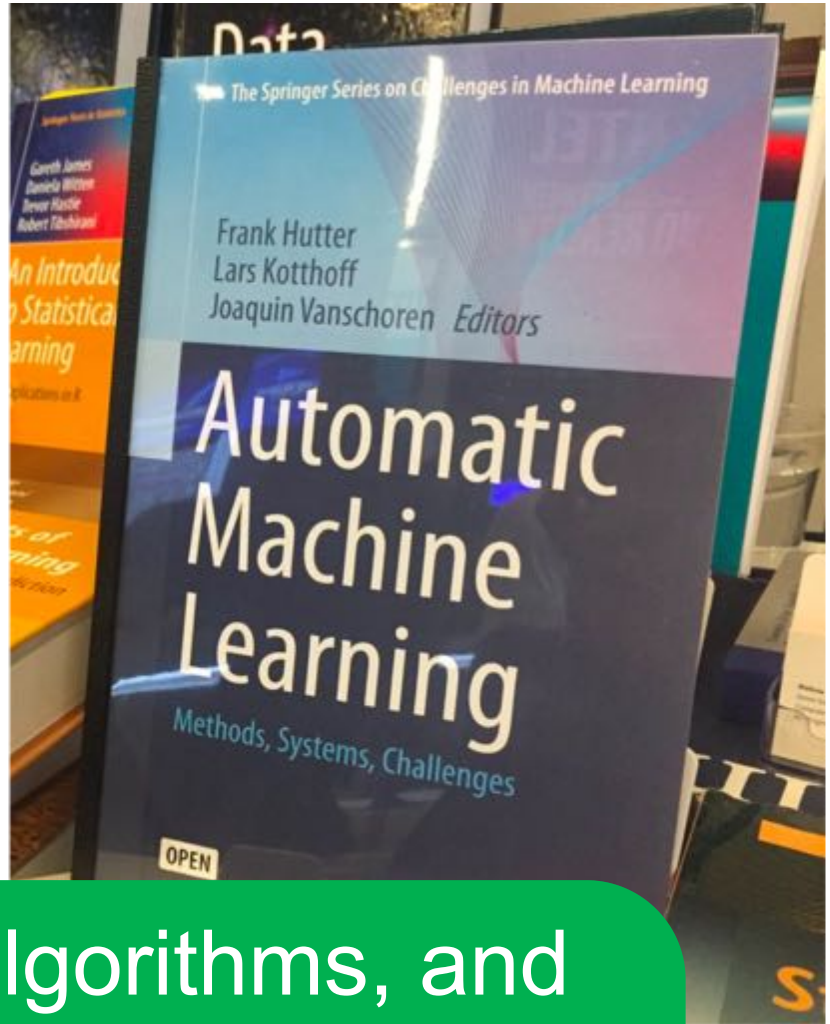

Provide the foundations, algorithms, and tools to develop systems that ease and support building ML/AI models as much as possible and in turn help reproducing and hopfeully even justifying our results

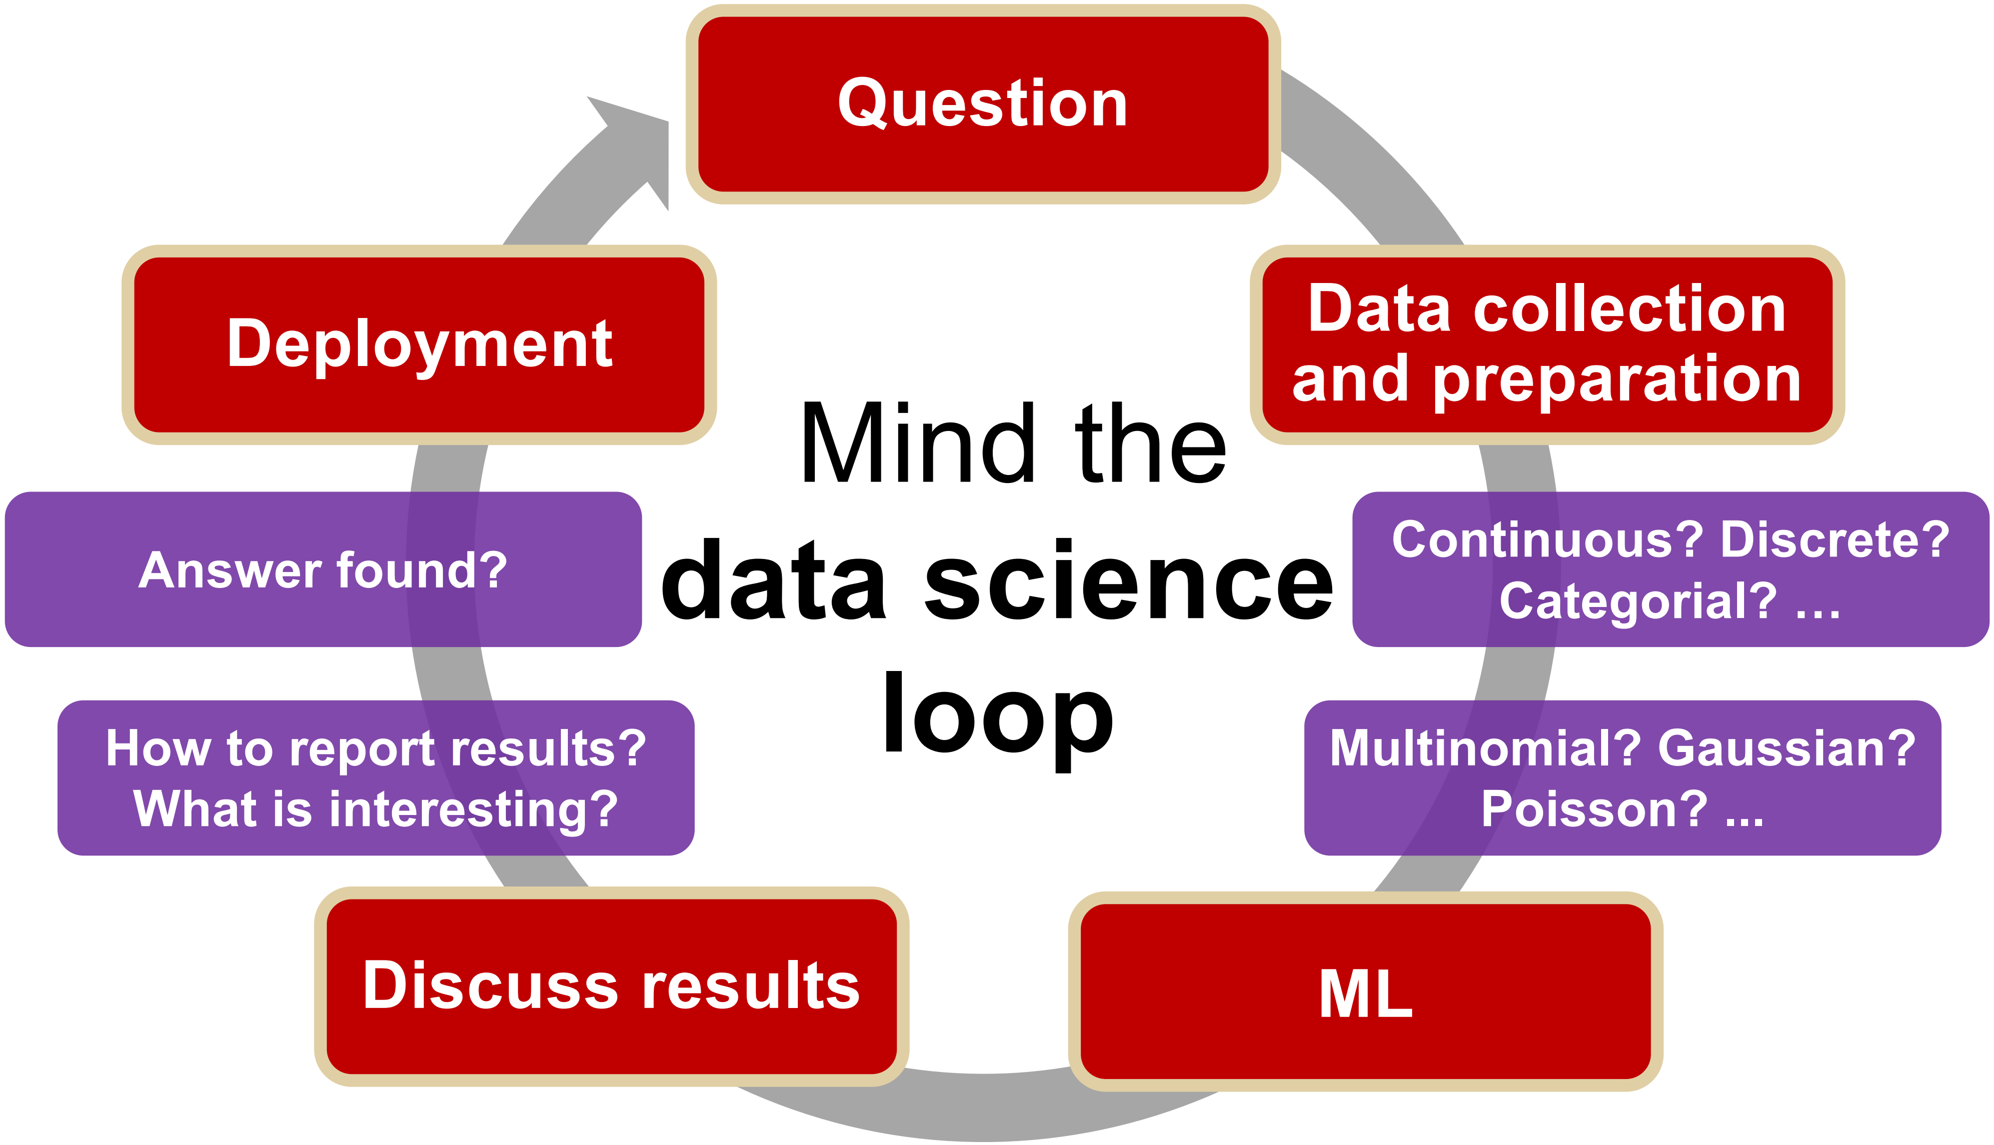
Question
Data collection and preparation
Continuous? Discrete? Categorial? …
Multinomial? Gaussian? Poisson? …
ML
Discuss results
How to report results? What is interesting?
Answer found?
Deployment
Mind the data science loop

[Vergari, Molina, Peharz, Ghahramani, Kersting, Valera AAAI 2019]

CAML DFG MADESI Federal Ministry of Education and Research

The Automatic Data Scientist

UBER AI Labs · UNIVERSITY OF CAMBRIDGE · Max Planck Institute for Intelligent Systems · TECHNISCHE UNIVERSITÄT DARMSTADT

AAAI-19: Thirty-Third AAAI Conference on Artificial Intelligence
January 27 – February 1, 2019, Hilton Hawaiian Village, Honolulu, Hawaii, USA

missing value

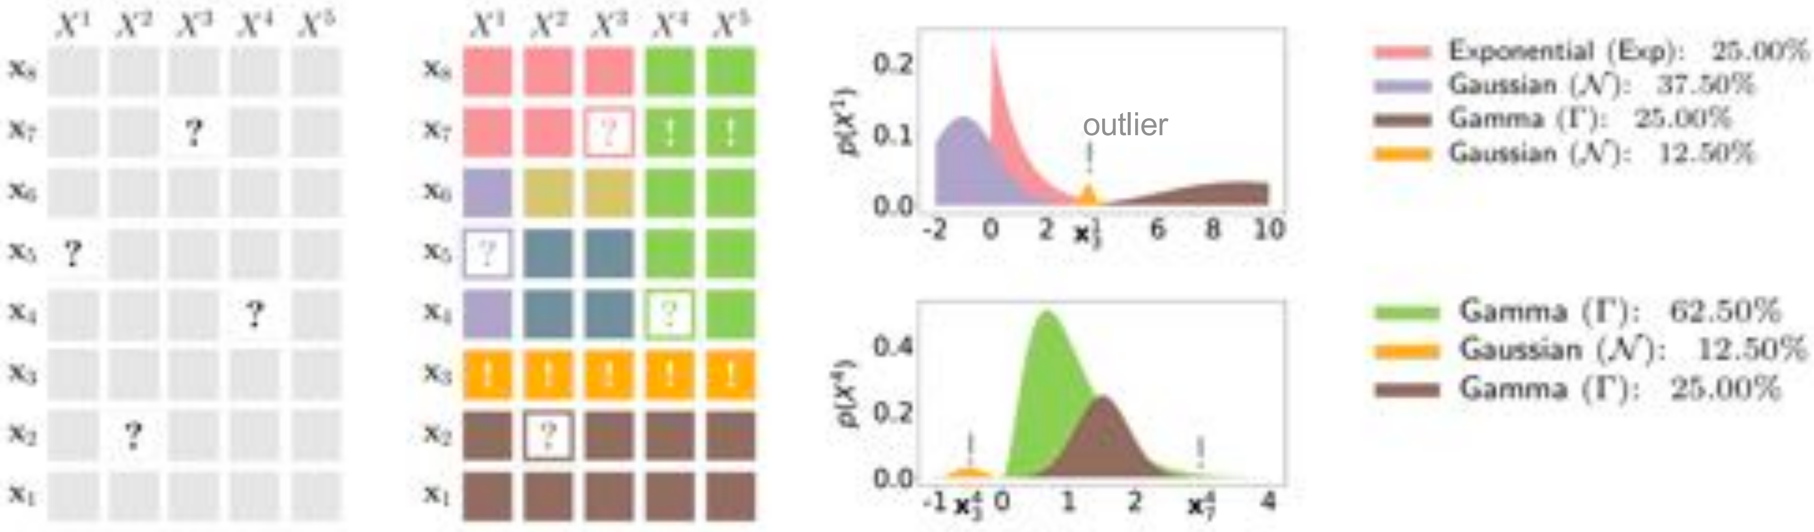

We can even automatically discovers the statistical types and parametric forms of the variables

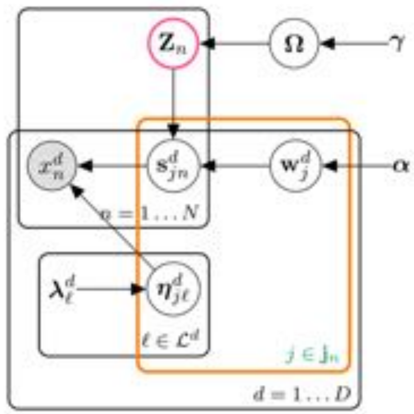

Bayesian Type Discovery

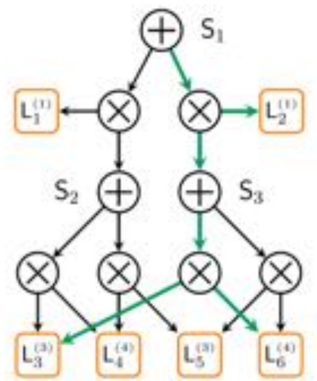

Mixed Sum-Product Network

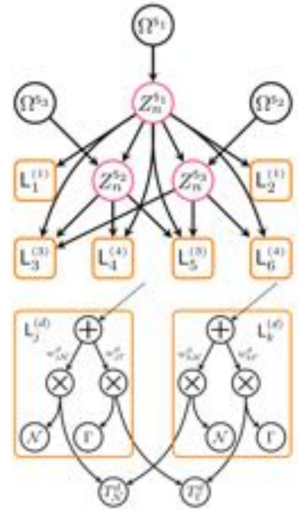

Automatic Statistician

That is, the machine understands the data with few expert input …

[Voelcker, Molina, Neumann, Westermann, Kersting ADS 2019]

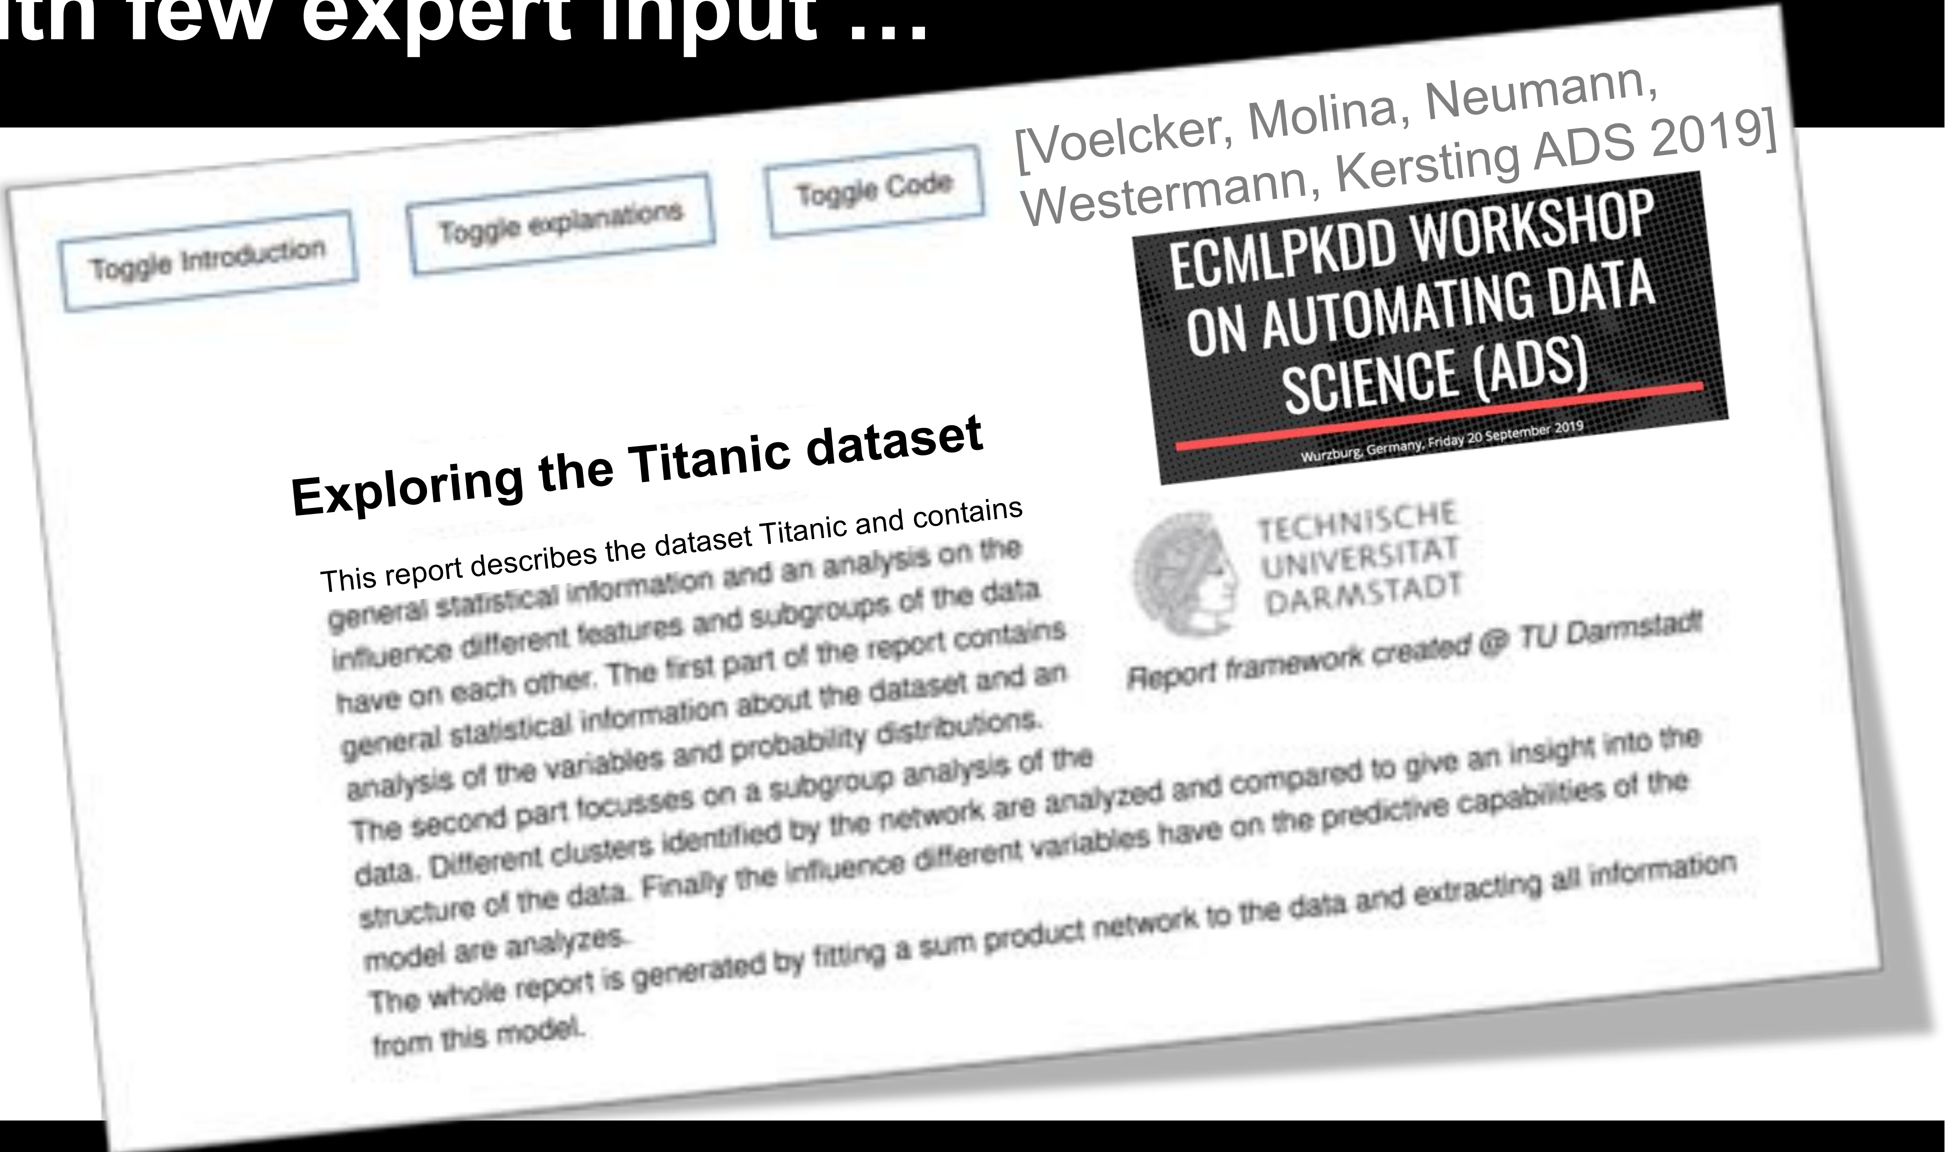
Toggle Introduction
Toggle explanations
Toggle Code

ECMLPKDD WORKSHOP ON AUTOMATING DATA SCIENCE (ADS)

Wurzburg, Germany, Friday 20 September 2019

TECHNISCHE UNIVERSITAT DARMSTADT

Report framework created @ TU Darmstadt

Exploring the Titanic dataset

This report describes the dataset Titanic and contains general statistical information and an analysis on the influence different features and subgroups of the data have on each other. The first part of the report contains general statistical information about the dataset and an analysis of the variables and probability distributions. The second part focusses on a subgroup analysis of the data. Different clusters identified by the network are analyzed and compared to give an insight into the structure of the data. Finally the influence different variables have on the predictive capabilities of the model are analyzes.
The whole report is generated by fitting a sum product network to the data and extracting all information from this model.

…and can compile data reports automatically

Programming languages for Systems AI,

the computational and mathematical modeling of complex AI systems.

[Laue et al. NeurIPS 2018; Kordjamshidi, Roth, Kersting: "Systems AI: A Declarative Learning Based Programming Perspective." IJCAI-ECAI 2018]

Th y

…bu e,

d

Eric Schmidt, Executive Chairman, Alphabet Inc.: Just Say "Yes", Stanford Graduate School of Business, May 2, 2017.https://www.youtube.com/watch?v=vbb-AjiXyh0.

Since science is more than a single table !

P(heart attack | )?

Crossover of ML and AI with data & programming abstractions

De Raedt, Kersting, Natarajan, Poole: Statistical Relational Artificial Intelligence: Logic, Probability, and Computation. Morgan and Claypool Publishers, ISBN: 9781627058414, 2016.

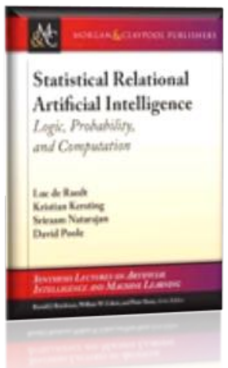

KATHOLIEKE UNIVERSITEIT LEUVEN

THE UNIVERSITY OF TEXAS AT DALLAS

TECHNISCHE UNIVERSITÄT DARMSTADT

UBC THE UNIVERSITY OF BRITISH COLUMBIA

building general-purpose AI and ML machines

make the ML/AI expert more effective

increases the number of people who can successfully build ML/AI applications

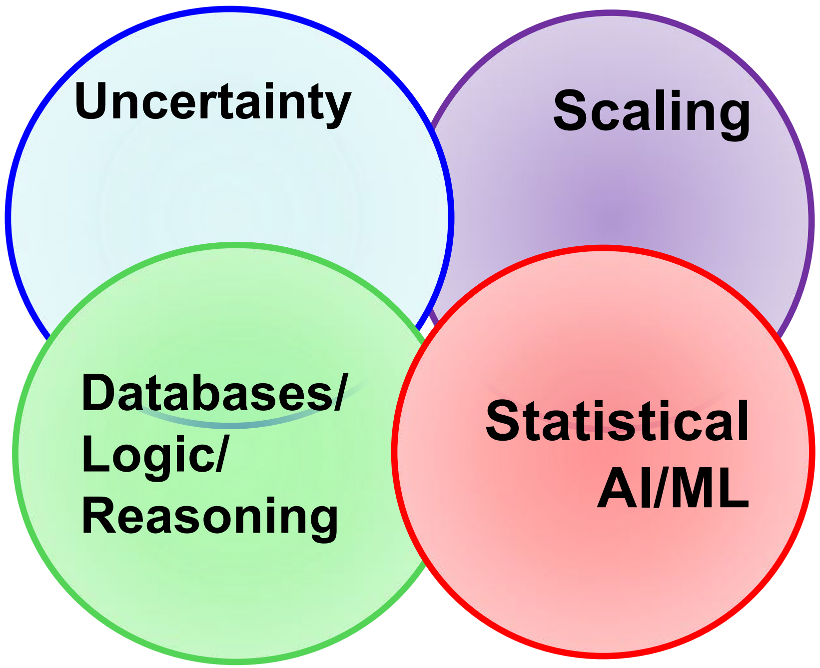

Natarajan, Khot, Kersting, Shavlik. Boosted Statistical Relational Learners. Springer Brief 2015

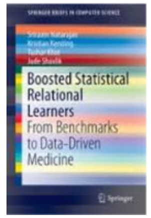

Understanding Electronic Health Records

Atherosclerosis is the cause of the majority of Acute Myocardial Infarctions (heart attacks)

TECHNISCHE UNIVERSITÄT DARMSTADT

THE UNIVERSITY OF TEXAS AT DALLAS

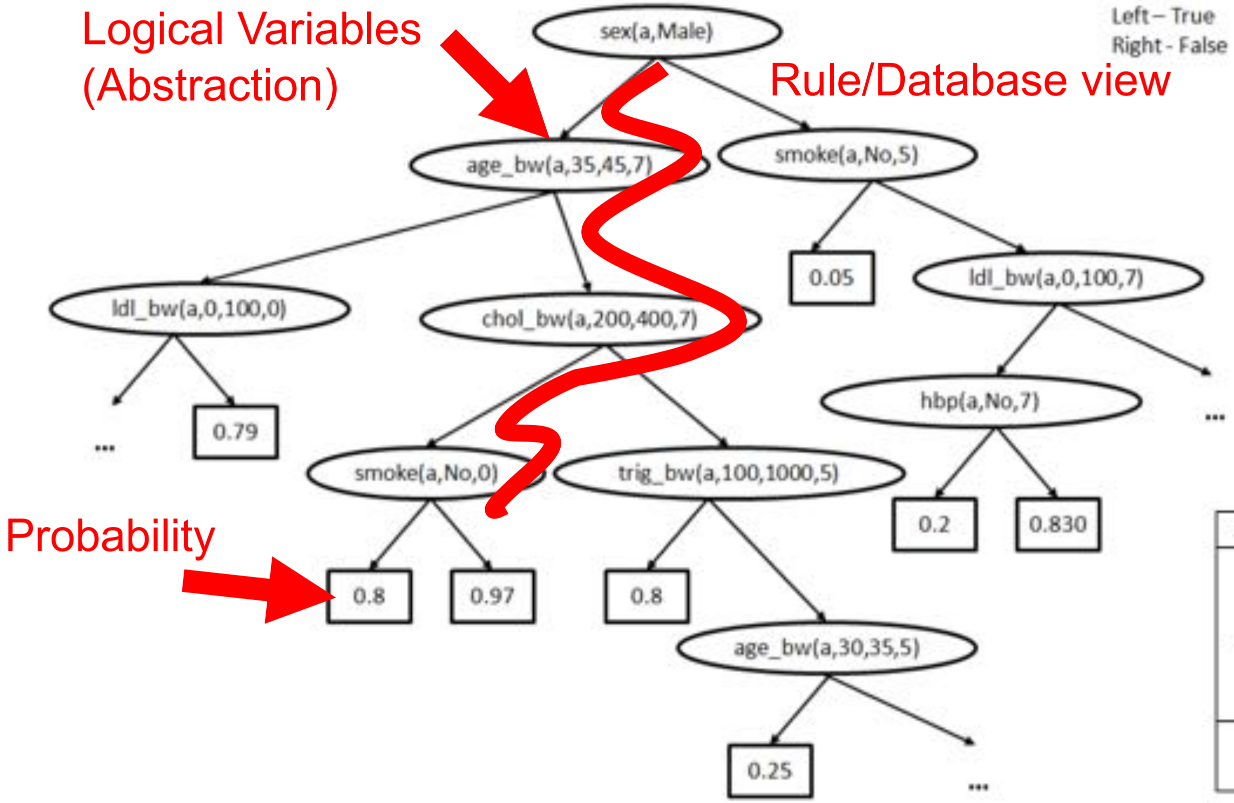

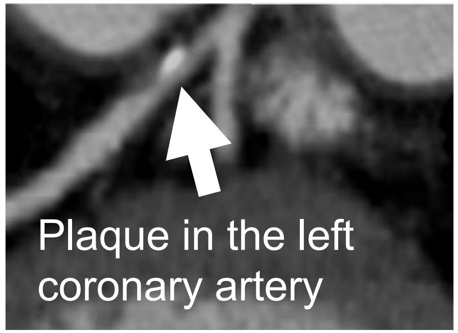

[Circulation; 92(8), 2157-62, 1995; JACC; 43, 842-7, 2004]

Algorithm	Accuracy	AUC-ROC
J48	0.667	0.607
SVM	0.667	0.5
AdaBoost	0.667	0.608
Bagging	0.677	0.613
NB	0.75	0.653
RPT	0.669*	0.778
RFGB	0.667*	0.819

The higher, the better

25%

Algorithm for Mining Markov Logic Networks	Likelihood The higher, the better	AUC-ROC The higher, the better	AUC-PR The higher, the better	Time The lower, the better
Boosting	0.81	0.96	0.93	**9s**
LSM	0.73	0.54	0.62	**93 hrs**

11% 78% 50% 37200x faster

state-of-the-art

[Kersting, Driessens ICML´08; Karwath, Kersting, Landwehr ICDM´08; Natarajan, Joshi, Tadepelli, Kersting, Shavlik. IJCAI´11; Natarajan, Kersting, Ip, Jacobs, Carr IAAI `13; Yang, Kersting, Terry, Carr, Natarajan AIME ´15; Khot, Natarajan, Kersting, Shavlik ICDM´13, MLJ´12, MLJ´15, Yang, Kersting, Natarajan BIBM`17]

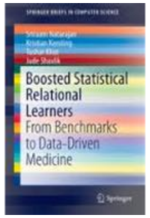

Natarajan, Khot, Kersting, Shavlik. Boosted Statistical Relational Learners. Springer Brief 2015

TECHNISCHE UNIVERSITÄT DARMSTADT

UTD THE UNIVERSITY OF TEXAS AT DALLAS

https://starling.utdallas.edu/software/boostsrl/wiki/

StARLinG LAB — People | Publications | Projects | Software | Datasets | Blog

BOOSTSRL BASICS

- Getting Started
- File Structure
- Basic Parameters
- Advanced Parameters
- Basic Modes
- Advanced Modes

ADVANCED BOOSTSRL

- Default (RDN-Boost)
- MLN-Boost
- Regression
- One-Class Classification
- Cost-Sensitive SRL
- Learning with Advice
- Approximate Counting
- Discretization of Continuous-Valued Attributes
- Lifted Relational Random Walks
- Grounded Relational Random Walks

APPLICATIONS

- Natural Language Processing

BoostSRL Wiki

BoostSRL (Boosting for Statistical Relational Learning) is a gradient-boosting based approach to learning different types of SRL models. As with the standard gradient-boosting approach, our approach turns the model learning problem to learning a sequence of regression models. The key difference to the standard approaches is that we learn relational regression models i.e., regression models that operate on relational data. We assume the data in a predicate logic format and the output are essentially first-order regression trees where the inner nodes contain conjunctions of logical predicates. For more details on the models and the algorithm, we refer to our book on this topic.

Sriraam Natarajan, Tushar Khot, Kristian Kersting and Jude Shavlik, Boosted Statistical Relational Learners: From Benchmarks to Data-Driven Medicine . SpringerBriefs in Computer Science, ISBN: 978-3-319-13643-1, 2015

Human-in-the-loop learning

Not every scientist likes to turn math into code

$$\min_{\mathbf{w},b,\boldsymbol{\xi}} \; \mathcal{P}(\mathbf{w},b,\boldsymbol{\xi}) = \frac{1}{2}\mathbf{w}^2 + C\sum_{i=1}^{n}\xi_i$$

$$\text{subject to} \; \begin{cases} \forall i \quad y_i(\mathbf{w}^\top \Phi(\mathbf{x}_i) + b) \geq 1 - \xi_i \\ \forall i \quad \xi_i \geq 0 \end{cases}$$

Support Vector Machines
Cortes, Vapnik MLJ 20(3):273-297, 1995

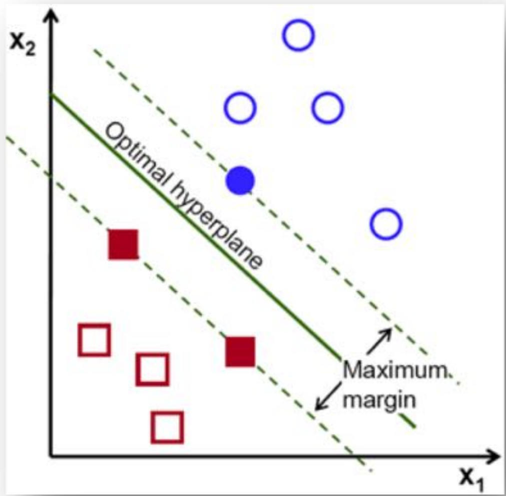

High-level Languages for Mathematical Programs

TECHNISCHE UNIVERSITÄT DARMSTADT

Write down SVM in „paper form.“ The machine compiles it into solver form.

```
#QUADRATIC OBJECTIVE
minimize: sum{J in feature(I,J)} weight(J)**2 + c1 * slack + c2 * coslack;

#labeled examples should be on the correct side
subject to forall {I in labeled(I)}: labeled(I)*predict(I) >= 1 - slack(I);

#slacks are positive
subject to forall {I in labeled(I)}: slack(I) >= 0;
```

Embedded within Python s.t. loops and rules can be used

reloop

RELOOP: A Toolkit for Relational Convex Optimization

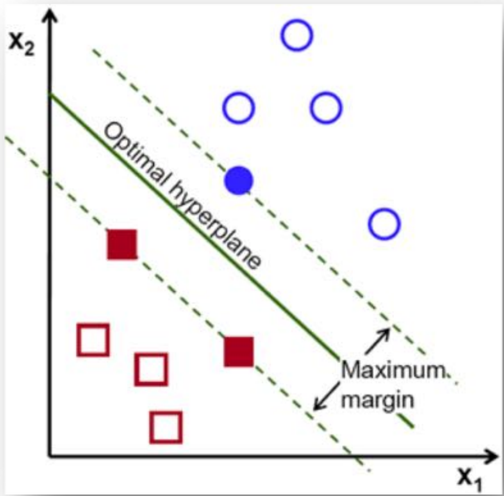

Support Vector Machines
Cortes, Vapnik MLJ 20(3):273-297, 1995

There are strong invests into high-level programming languages for AI/ML

RelationalAI, Apple, Microsoft and Uber are investing hundreds of millions of US dollars

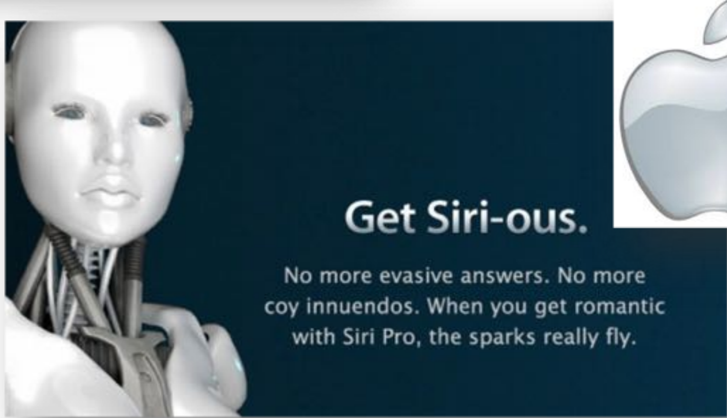

UNIVERSITY OF CAMBRIDGE

UBER AI Labs

Max Planck Institute for Intelligent Systems

UNI GRAZ

FLOW

TECHNISCHE UNIVERSITÄT DARMSTADT

Microsoft Research

UBER AI Labs

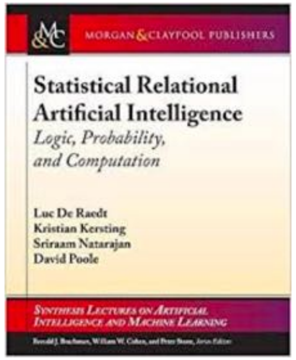

relationalAI

AI for the enterprise

Getting deep systems that reason and know what they don't know

Responsible AI systems that explain their decisions and co-evolve with the humans

Open AI systems that are easy to realize and understandable for the domain experts

„Tell the AI when it is right for the wrong reasons and it adapts ist behavior"

(a) Original Image (b) Explaining *Electric guitar* (c) Explaining *Acoustic guitar* (d) Explaining *Labrador*

Figure 4: Explaining an image classification prediction made by Google's Inception network, highlighting positive pixels. The top 3 classes predicted are "Electric Guitar" ($p = 0.32$), "Acoustic guitar" ($p = 0.24$) and "Labrador" ($p = 0.21$)

Teso, Kersting AIES 2019

AAAI / ACM conference on ARTIFICIAL INTELLIGENCE, ETHICS, AND SOCIETY

Making Clever Hans Clever

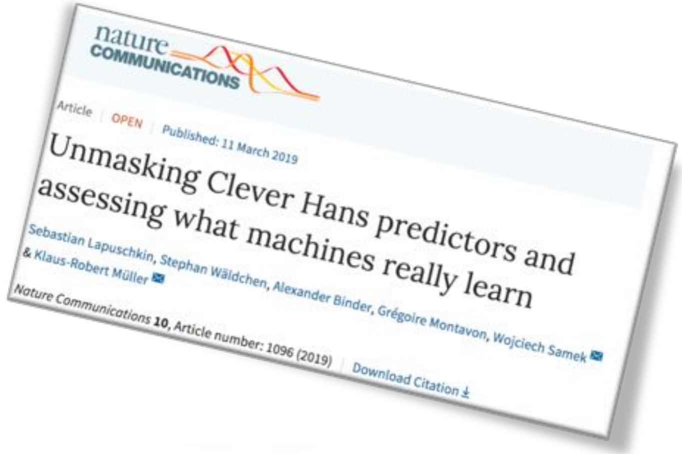
nature COMMUNICATIONS

Article | OPEN | Published: 11 March 2019

Unmasking Clever Hans predictors and assessing what machines really learn

Sebastian Lapuschkin, Stephan Wäldchen, Alexander Binder, Grégoire Montavon, Wojciech Samek & Klaus-Robert Müller

Nature Communications 10, Article number: 1096 (2019) | Download Citation

Co-adaptive ML:

- human is changing computer behavior
- human adapts his or her data and goals in response to what is learned

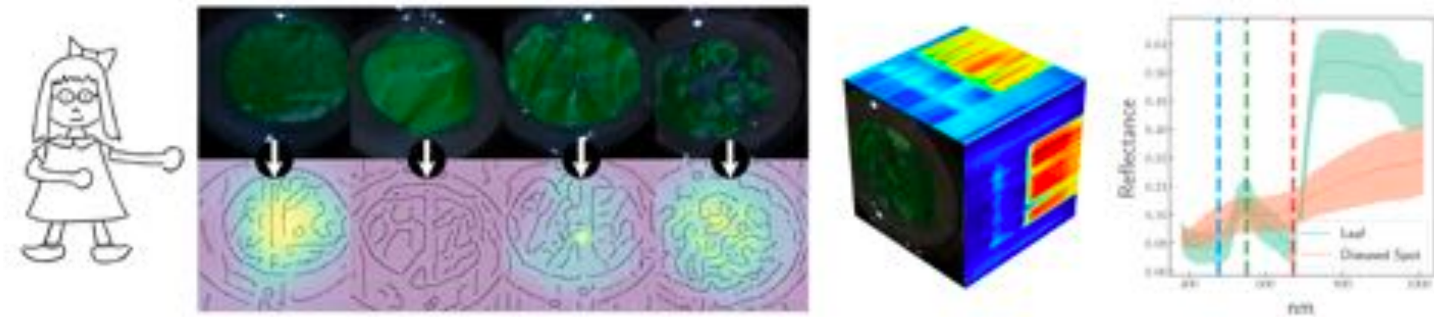

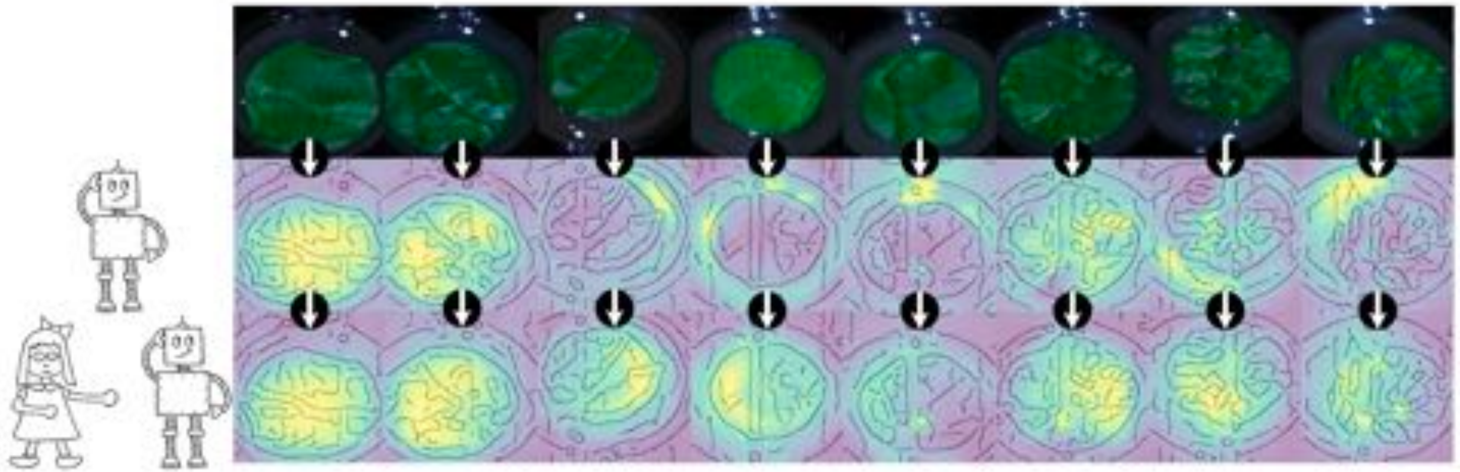

[Teso, Kersting AIES 2019 and ongoing research]

AAAI / ACM conference on ARTIFICIAL INTELLIGENCE, ETHICS, AND SOCIETY

Explanation should be understandable by humans

The twin science: cognitive science

"How do we humans get so much from so little?" and by that I mean how do we acquire our understanding of the world given what is clearly by today's engineering standards so little data, so little time, and so little energy.

Centre for Cognitive Science at TU Darmstadt

Establishing cognitive science at the Technische Universität Darmstadt is a long-term commitment across multiple departments (see Members to get an impression on the interdisciplinary of the supporting groups and departments). The TU offers a strong foundation including several established top engineering groups in Germany, a prominent computer science department (which is among the top four in Germany), a

Centre for Cognitive Science

Josh Tenenbaum, MIT

Lake, Salakhutdinov, Tenenbaum, Science 350 (6266), 1332-1338, 2015

Tenenbaum, Kemp, Griffiths, Goodman, Science 331 (6022), 1279-1285, 2011

The future of AI
The third wave of AI

Data are now ubiquitous; there is great value from understanding this data, building models and making predictions

However, data is not everything

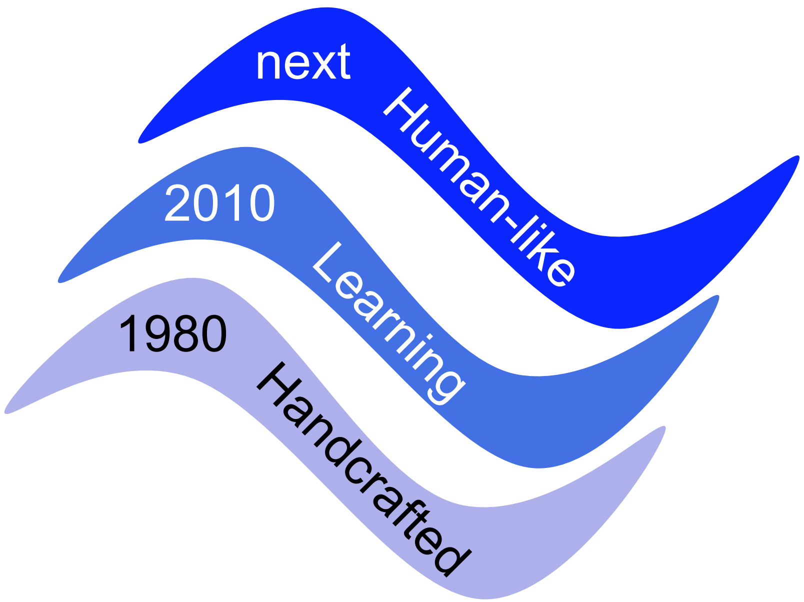

AI systems that can acquire human-like communication and reasoning capabilities, with the ability to recognise new situations and adapt to them.

Meeting this grand challenge is a team sport !

And this is AI! Still a lot to be done! It is a team sport.

Thanks to all students of the Research Training Group "Artificial Intelligence - Facts, Chances, Risks" of the German National Academic Scholarship Foundation. Special thanks to **Maike Elisa Müller** and **Jannik Kossen** for taking the lead and to **Matthias Kleiner**, president of the Leibniz Association, for his preface

Illustration Nanina Föhr