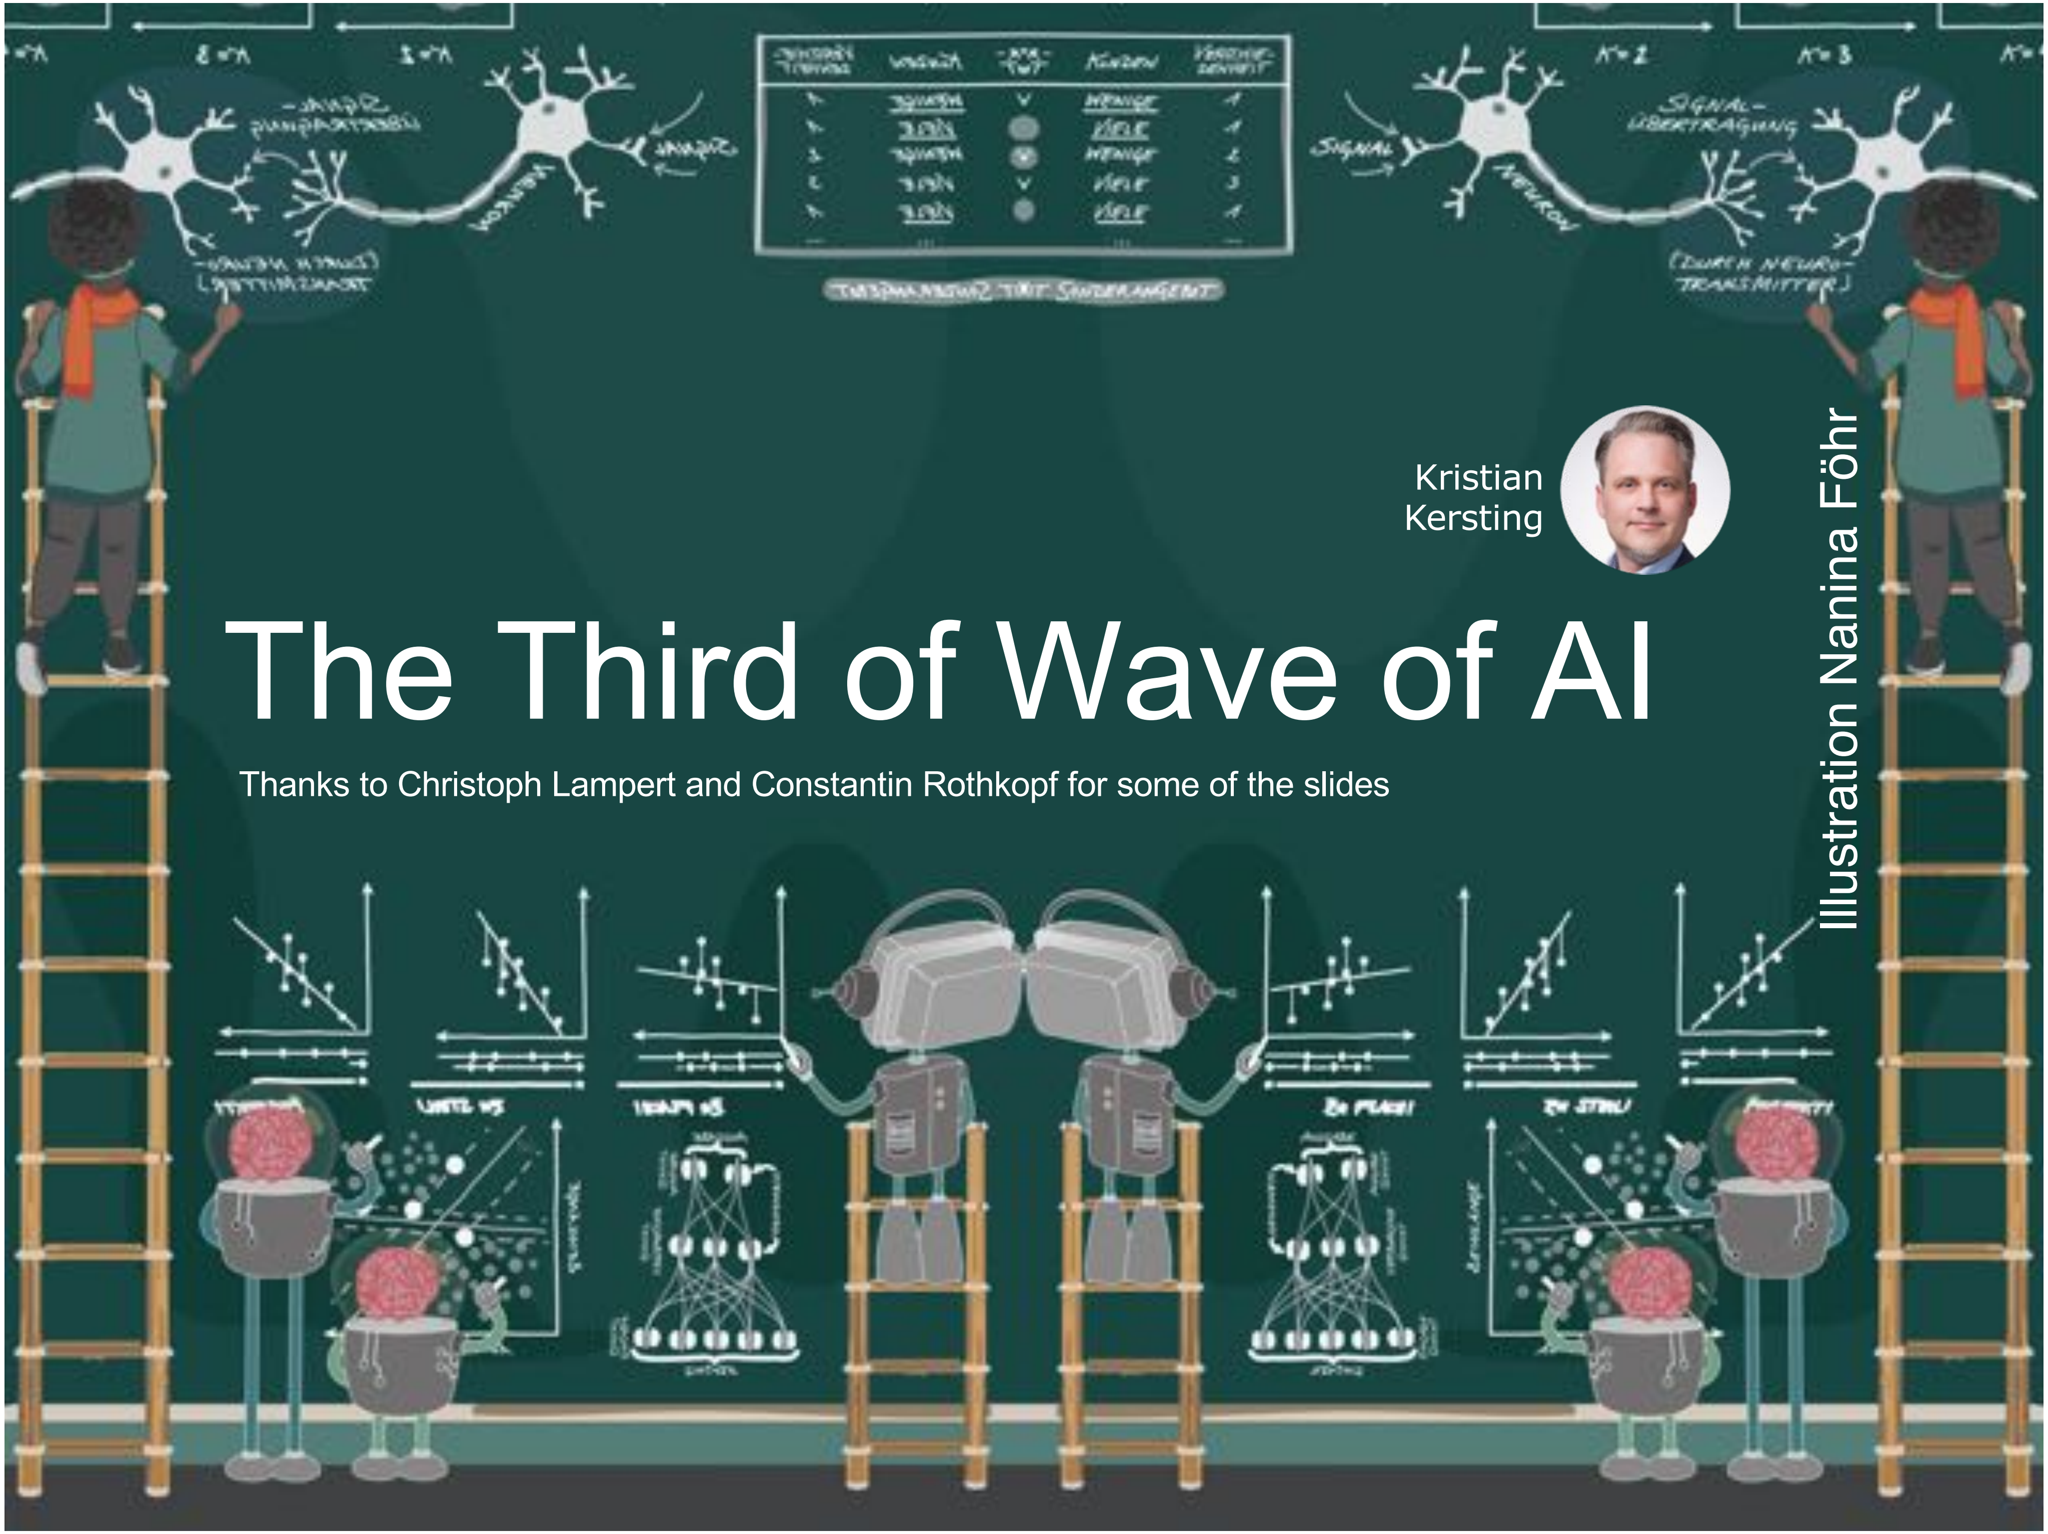

The Third of Wave of AI

Thanks to Christoph Lampert and Constantin Rothkopf for some of the slides

Kristian
Kersting

Illustration Nanina Föhr

The dream of an artificially intelligent entity is not new

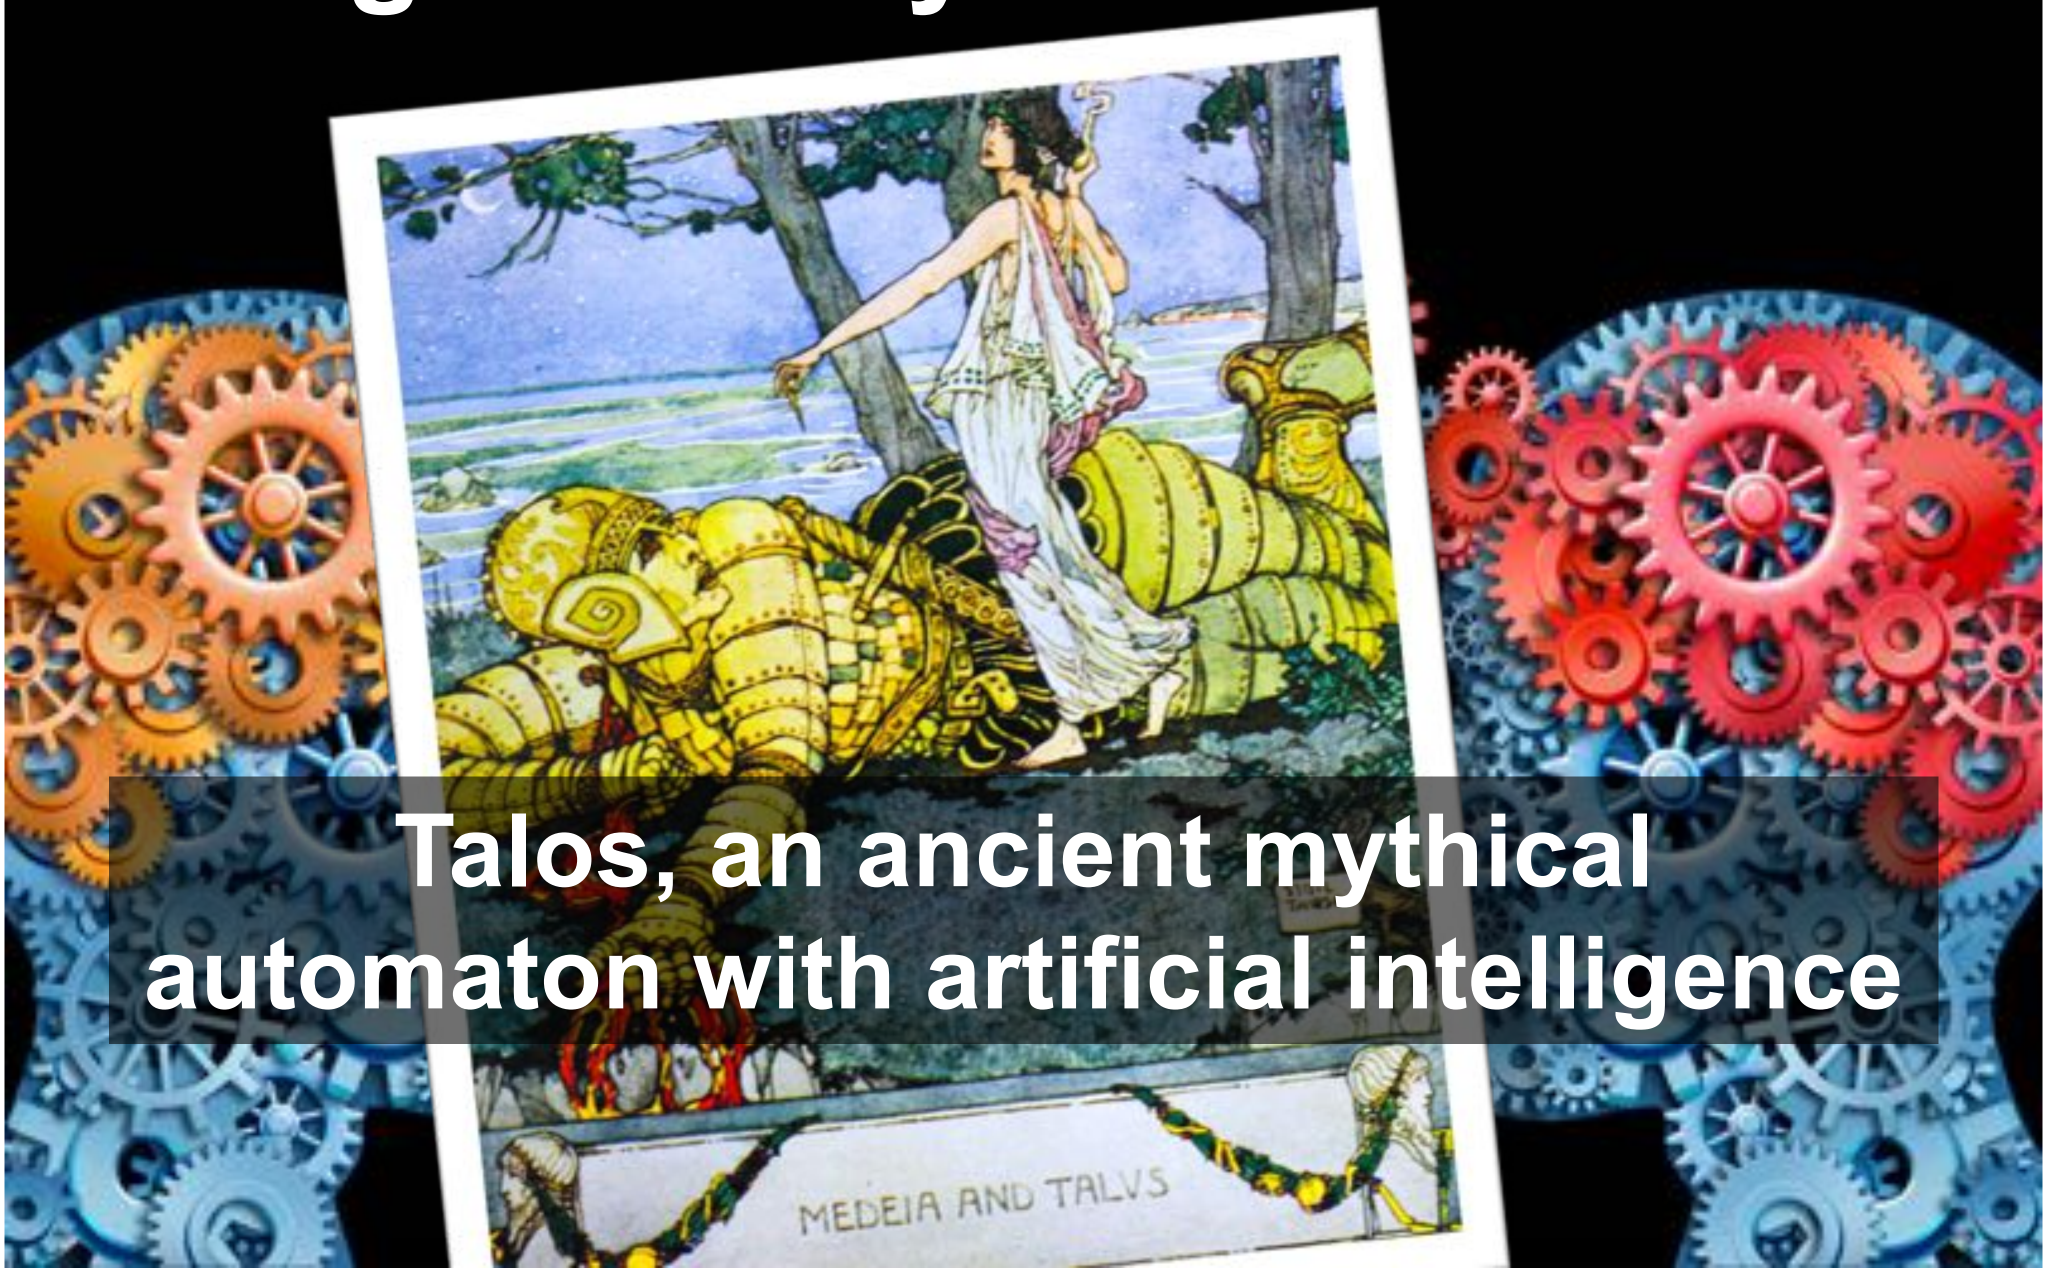

Talos, an ancient mythical automaton with artificial intelligence

The dream of an artificially intelligent entity is not new

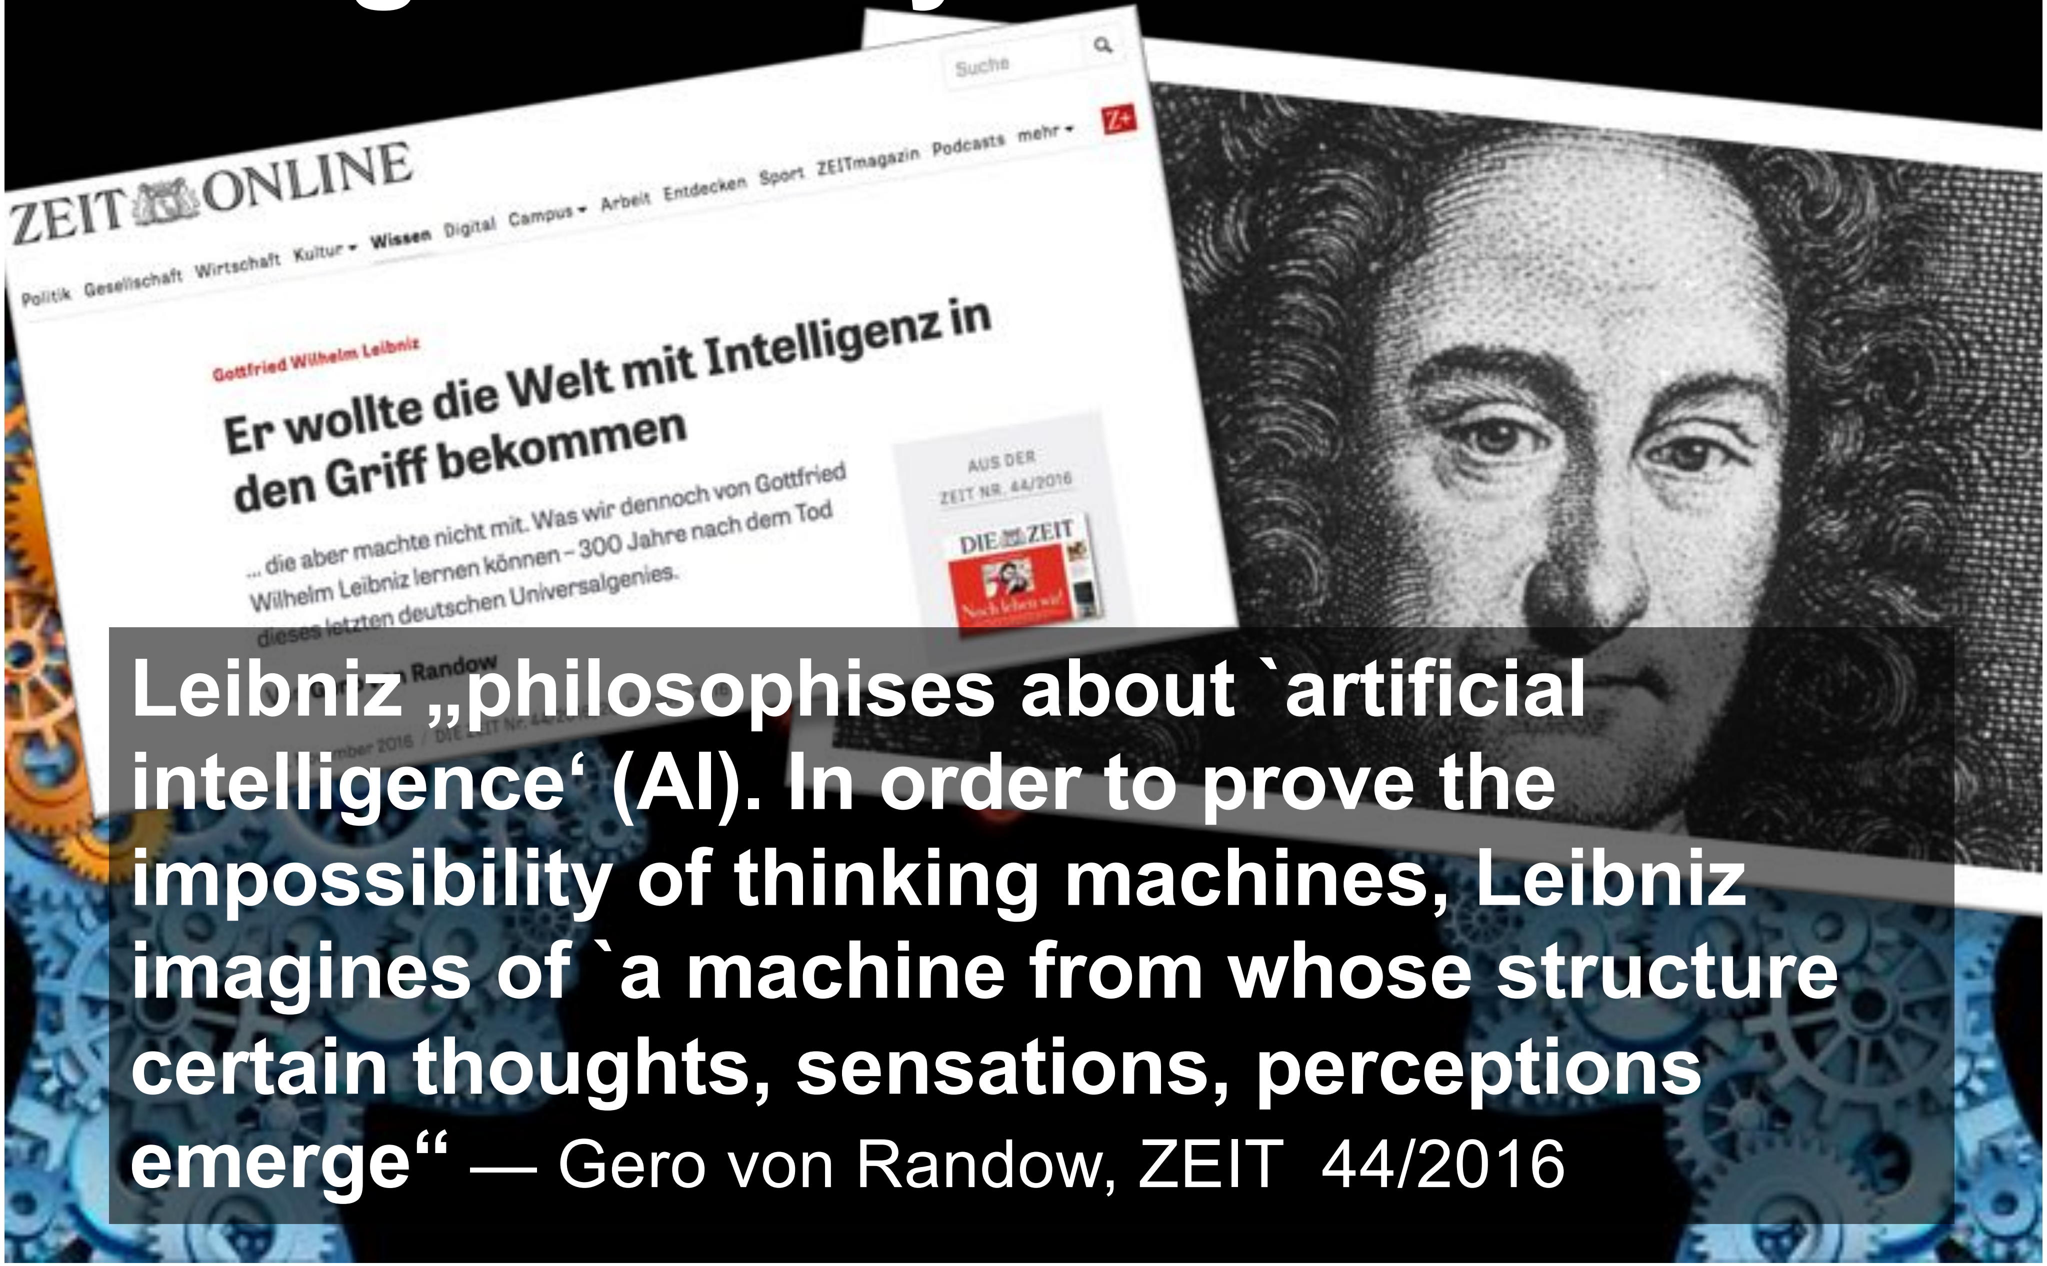

Leibniz „philosophises about `artificial intelligence‘ (AI). In order to prove the impossibility of thinking machines, Leibniz imagines of `a machine from whose structure certain thoughts, sensations, perceptions emerge“ — Gero von Randow, ZEIT 44/2016

AI today

the INQUIRER
Artificial intelligence will create the next industrial revolution, experts claim
Artificial intelligence better than scientists at choosing successful embryos
'We won't waste time on treatments that won't work, so the patient should get
says clinic director
BBC NEWS
Technology
Stephen Hawking warns artificial intelligence could end mankind
"Humans, who are limited by slow biological evolution, couldn't compete and would be
Telegraph
HOME
Lifestyle · Cars · News
Self-driving Tesla 'sav
by steering him to hos
V/S
Elon Musk
@elonmusk
I've talked to Mark about this. His understanding of the subject is limited.
SCIENTIFIC AMERICAN DECEMBER 2016
Computers Now Recognize Patterns Better Than Humans Can
An approach to artificial intelligence that enables computers to recognize visual patterns better than humans are able to do

AI today

THE ECONOMIC IMPACT OF
ARTIFICIAL INTELLIGENCE
NORTH AMERICA
$3.7 TR
NORTHERN EUROPE
$1.8 TR
CHINA
$7 TR
$0.7 TR
SOUTHERN EUROPE
LATIN AMERICA
$0.5 TR
$0.9 TR
DEVELOPED ASIA
REST OF WORLD
$1.2 TR
Projected Global Economic Effects of AI by 2030
Source: PwC

AI Impact driven by researchers

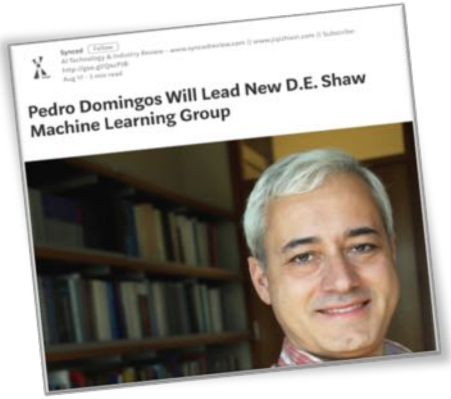

Pedro Domginos
UW, **DE Shaw**

Charles Elkan
UCSD, **Goldman Sachs**

UBER AI Labs

Zoubin Ghahramani
Cambridge, **Uber**

Geoffrey Hinton
Turing Awardee
DeepMind, U. Toronto
Vector Institute

VECTOR INSTITUTE

Yann LeCun
Turing Awardee
NYU, **Facebook**

Manuela Veloso
Former AAAI President
CMU, **JPMorgan**

Yoshua Bengio
Turing Awardee
Element.AI
Univ. Montreal

E AI

… and many more examples

So, AI has many faces

But, what exactly is AI?

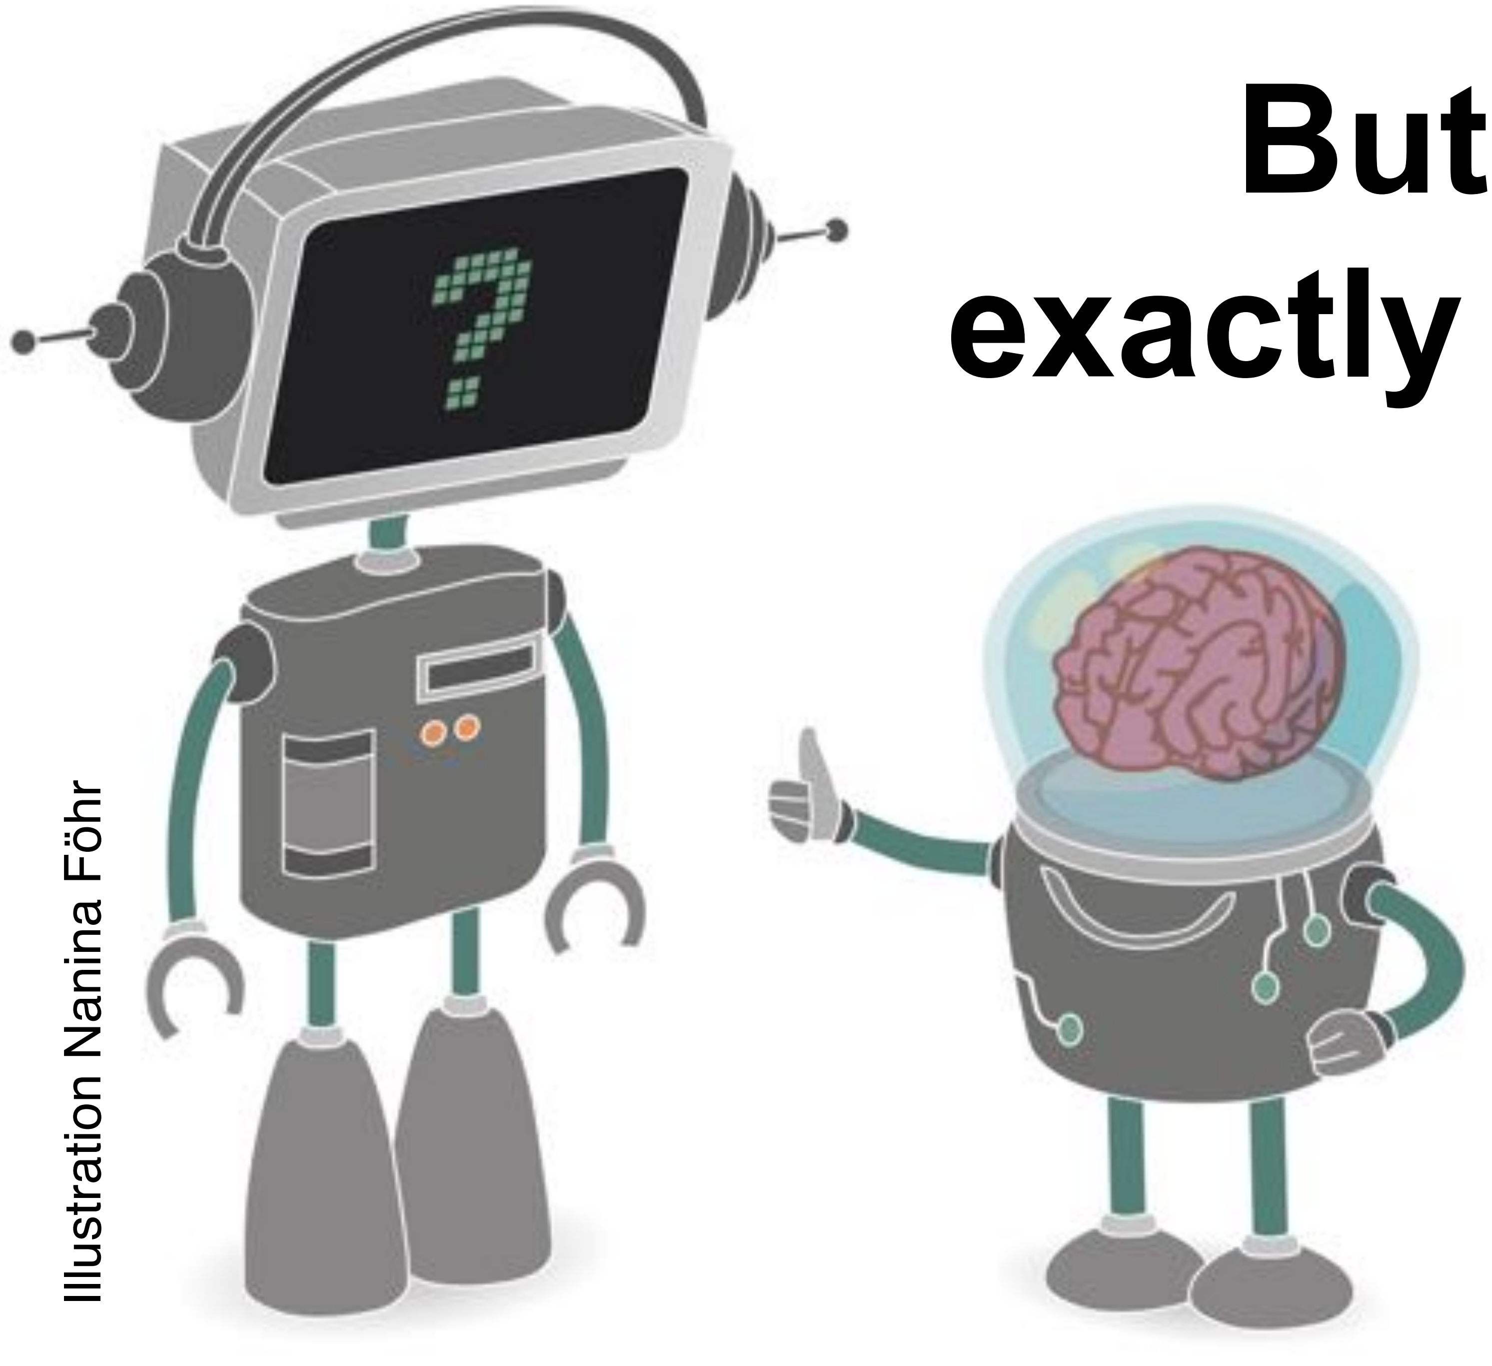

Illustration Nanina Föhr

Humans are considered to be smart

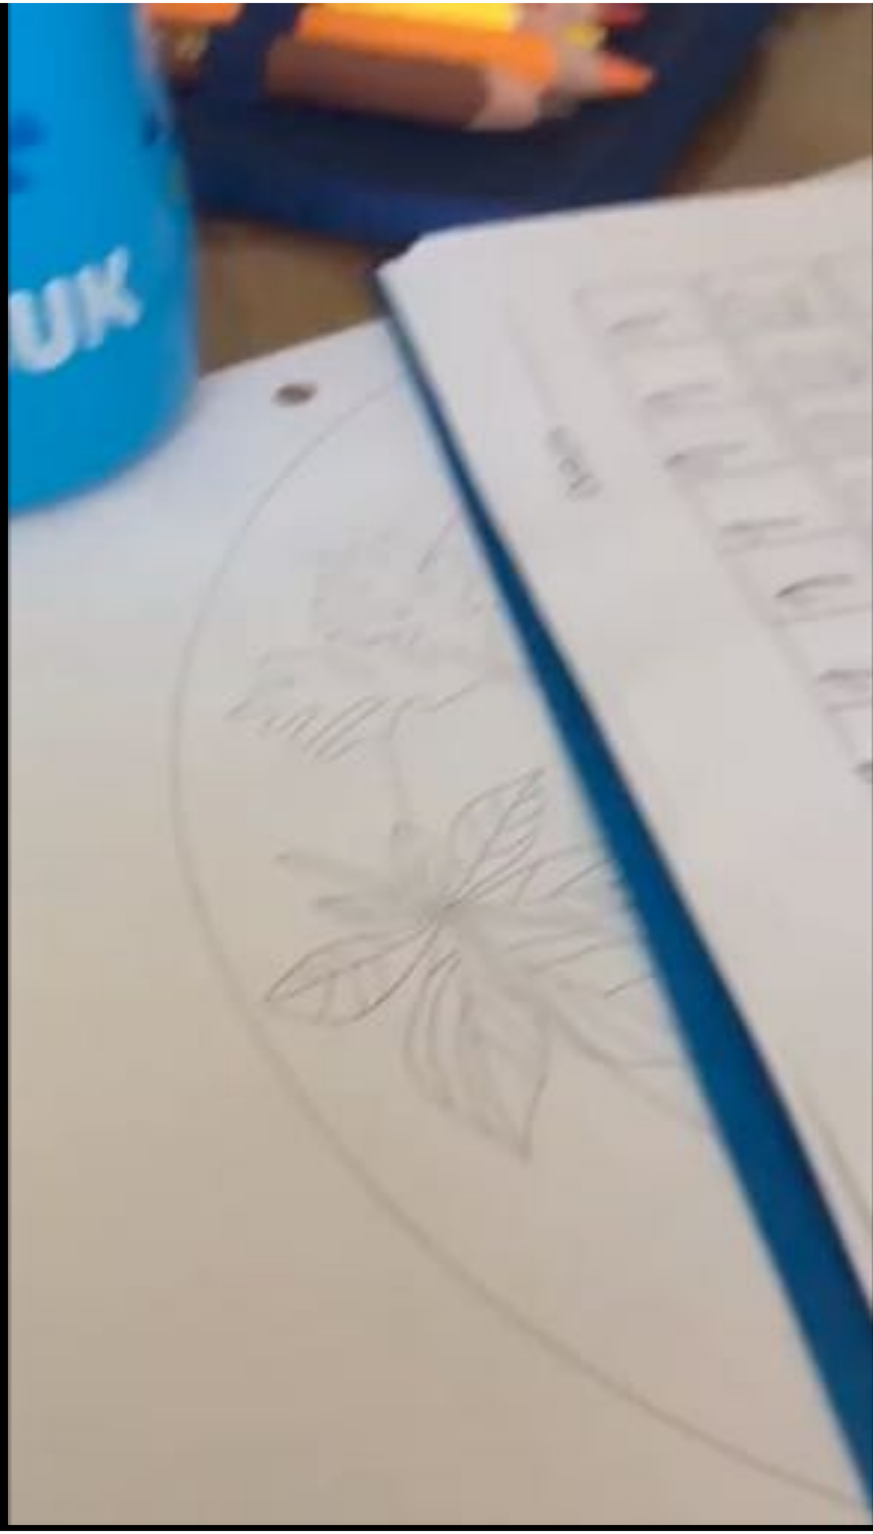

https://www.youtube.com/watch?v=XQ79UUIOeWc

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

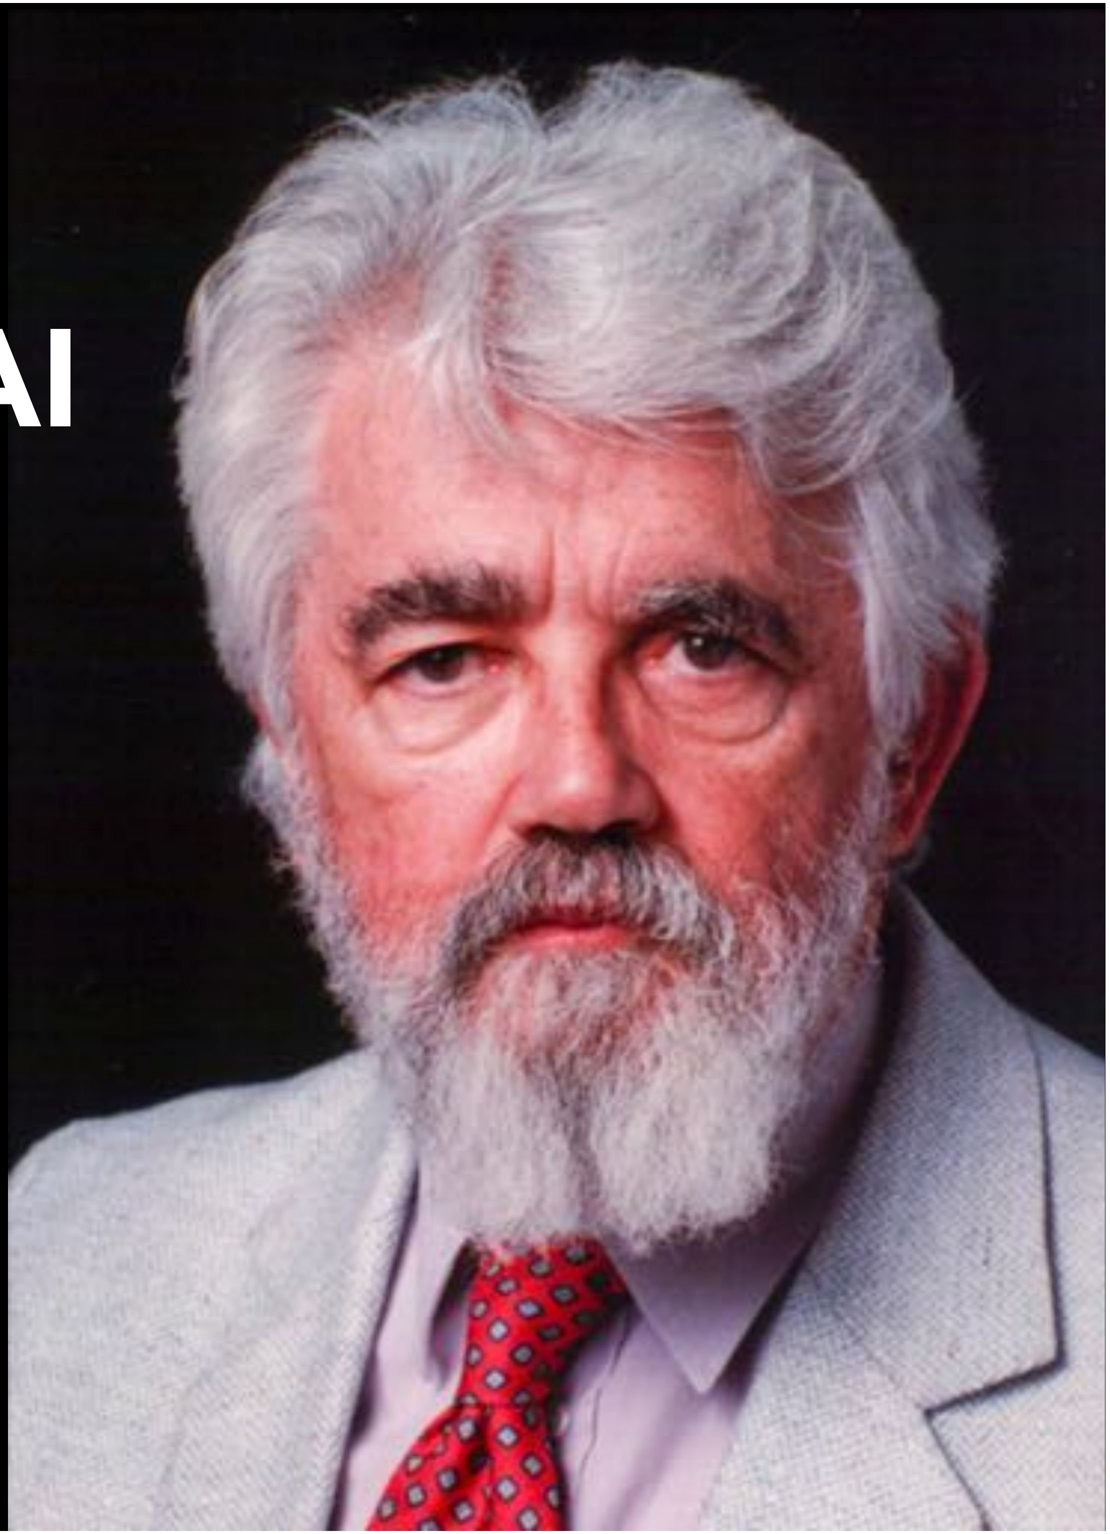

Learning Thinking Planning

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

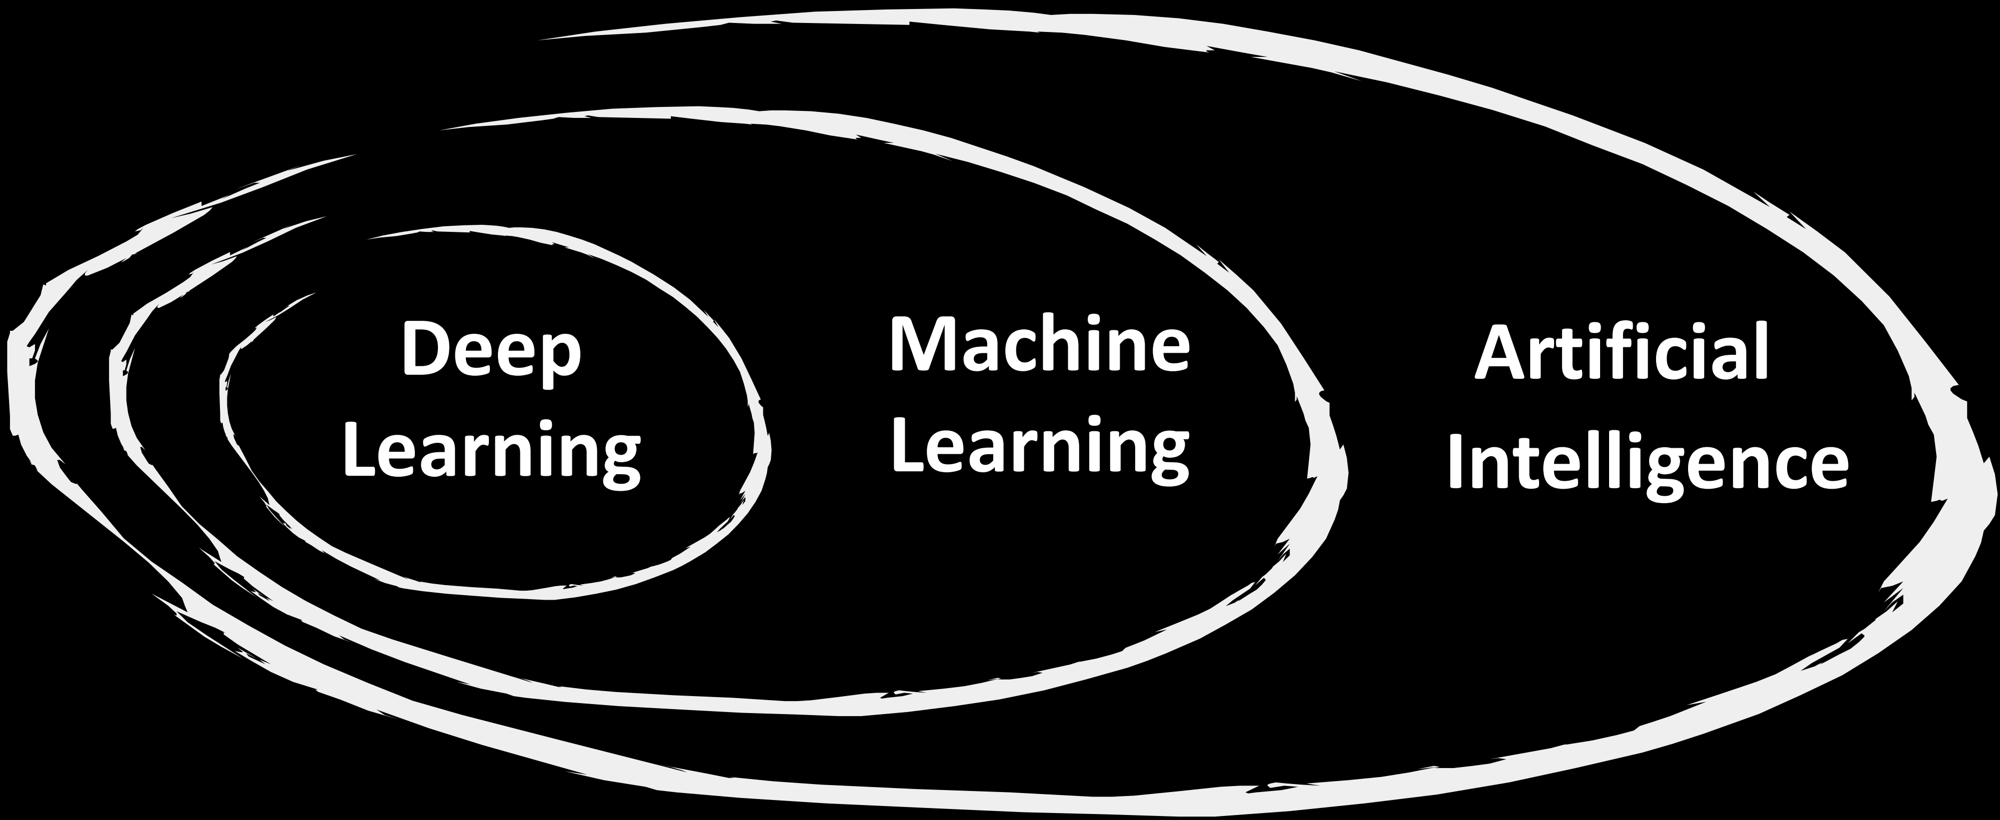

Three levels of description

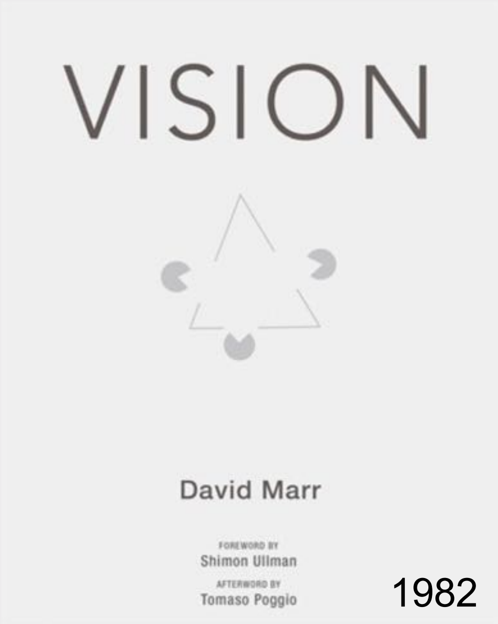

1982

Computational
Why do things work the way they work? What is the goal of the computation? What are the unifying principles?

maximize:

$$R_t = r_{t+1} + r_{t+2} + \cdots + r_T$$

Algorithmic
What represetation can implement such computations? How does the choice of the representation determine the algorithm

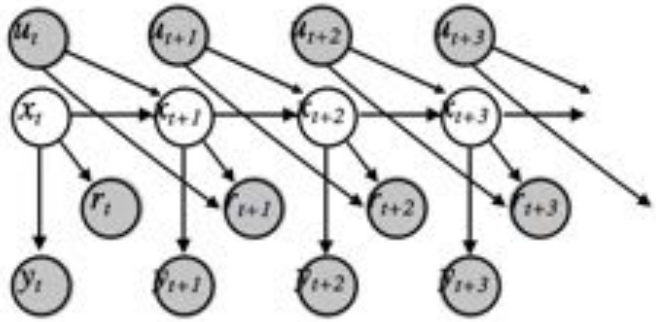

Implementational
How can such a system be built in hardware?
How can neurons carry out the computations?

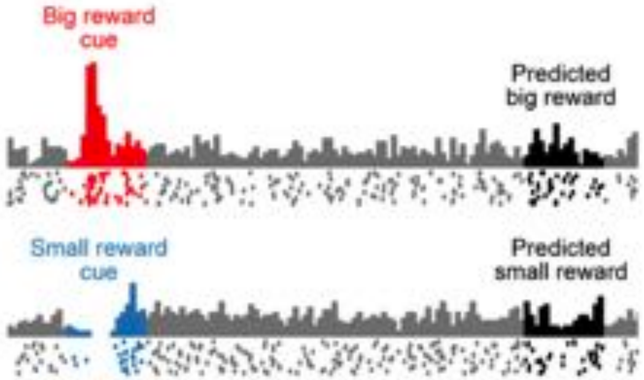

slide after C. Rothkopf (TUD)

And this all started as early as 1956

ONCE UPON A TIME

1956 Birth of AI

A Proposal for the

DARTMOUTH SUMMER RESEARCH PROJECT ON ARTIFICIAL INTELLIGENCE

We propose that a 2 month, 10 man study of artificial intelligence be carried out during the summer of 1956 at Dartmouth College in Hanover, New Hampshire. The study is to proceed on the basis of the conjecture that every aspect of learning or any other feature of intelligence can in principle be so precisely described that a machine can be made to simulate it. An attempt will be made to find how to make machines use language, form abstractions and concepts, solve kinds of problems now reserved for humans, and improve themselves. We think that a significant advance can be made in one or more of these problems if a carefully selected group of scientists work on it together for a summer.

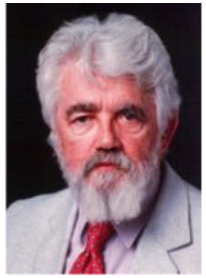

John McCarthy
Turing Award 1971

Marvin Minsky
Turing Award 1969

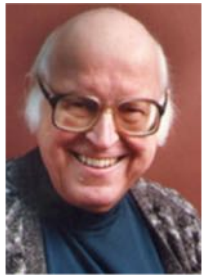

Allen Newell
Turing Award 1975

Herbert A. Simon
Turing Award 1975
Nobel Prize 1978

… and of Cognitive Science

Artificial Neural Networks

COGNITIVE SCIENCE **14**, 179–211 (1990)

Finding Structure in Time

JEFFREY L. ELMAN
University of California, San Diego

Learning representations by back-propagating errors

David E. Rumelhart*, Geoffrey E. Hinton† & Ronald J. Williams*

* Institute for Cognitive Science, C-015, University of California, San Diego, La Jolla, California 92093, USA
† Department of Computer Science, Carnegie-Mellon University, Pittsburgh, Philadelphia 15213, USA

COGNITIVE SCIENCE 9, 147–169 (1985)

A Learning Algorithm for Boltzmann Machines*

DAVID H. ACKLEY
GEOFFREY E. HINTON
Computer Science Department
Carnegie-Mellon University
TERRENCE J. SEJNOWSKI
Biophysics Department
The Johns Hopkins University

Biological Cybernetics

Biol. Cybernetics 36, 193–202 (1980)

Neocognitron: A Self-organizing Neural Network Model for a Mechanism of Pattern Recognition Unaffected by Shift in Position

Kunihiko Fukushima
NHK Broadcasting Science Research Laboratories, Kinuta, Setagaya, Tokyo, Japan

Psychological Review
1981, Vol. 88, No. 2, 135–170

0033-295X/81/8802-0135$00.75

Toward a Modern Theory of Adaptive Networks: Expectation and Prediction

Richard S. Sutton and Andrew G. Barto
Computer and Information Science Department
University of Massachusetts—Amherst

Psychological Review
Vol. 65, No. 6, 1958

THE PERCEPTRON: A PROBABILISTIC MODEL FOR INFORMATION STORAGE AND ORGANIZATION IN THE BRAIN[1]

F. ROSENBLATT
Cornell Aeronautical Laboratory

Artificial Neural Networks

COGNITIVE SCIENCE 4, 179–211 (1990)

Finding Structure in Time

JEFFREY L. ELMAN

University of California, San Diego

Learning representations by back-propagating errors

David E. Rumelhart*, Geoffrey E. Hinton† & Ronald J. Williams*

* Institute for Cognitive Science C-015, University of California, San Diego, La Jolla, California 92093, USA
† Department of Computer Science, Carnegie-Mellon University, Pittsburgh, Philadelphia 15213, USA

COGNITIVE SCIENCE, 147–169 (1985)

A Learning Algorithm for Boltzmann Machines*

DAVID H. ACKLEY
GEOFFREY E. HINTON
Computer Science Department
Carnegie-Mellon University
TERRENCE J. SEJNOWSKI
Biophysics Department
The Johns Hopkins University

Biological Cybernetics

Biol. Cybernetics 36, 193–202 (1980)

Neocognitron: A Self-organizing Neural Network Model for a Mechanism of Pattern Recognition Unaffected by Shift in Position

Kunihiko Fukushima

NHK Broadcasting Science Research Laboratories, Kinuta, Setagaya, Tokyo, Japan

Psychological Review
1981, Vol. 88, No. 2, 135–170

Copyright 1981 by the American Psychological Association, Inc.
0033-295X/81/8802-0135$00.75

Toward a Modern Theory of Adaptive Networks: Expectation and Prediction

Richard S. Sutton and Andrew G. Barto
Computer and Information Science Department
University of Massachusetts—Amherst

Psychological Review
Vol. 65, No. 6, 1958

THE PERCEPTRON: A PROBABILISTIC MODEL FOR INFORMATION STORAGE AND ORGANIZATION IN THE BRAIN[1]

F. ROSENBLATT

Cornell Aeronautical Laboratory

slide after C. Rothkopf (TUD), after J.Tenenbaum (MIT)

Algorithms of intelligent behaviour teach us a lot about ourselves

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.

Josh Tenenbaum, MIT

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

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

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

Well-established scientific discipline with international societies, selective venues, and networks

Association for the Advancement of Artificial Intelligence

IJCAI

ACL

HCI INTERNATIONAL CONFERENCE human-computer interaction

ICML

NeurIPS
Thirty-second Conference on Neural Information Processing Systems

CVPR

ISWC

ROBOTICS SCIENCE AND SYSTEMS

EurAi European Association for Artificial Intelligence

ellis European Laboratory for Learning and Intelligent Systems

CLAIRE CONFEDERATION OF LABORATORIES FOR ARTIFICIAL INTELLIGENCE RESEARCH IN EUROPE

… among others

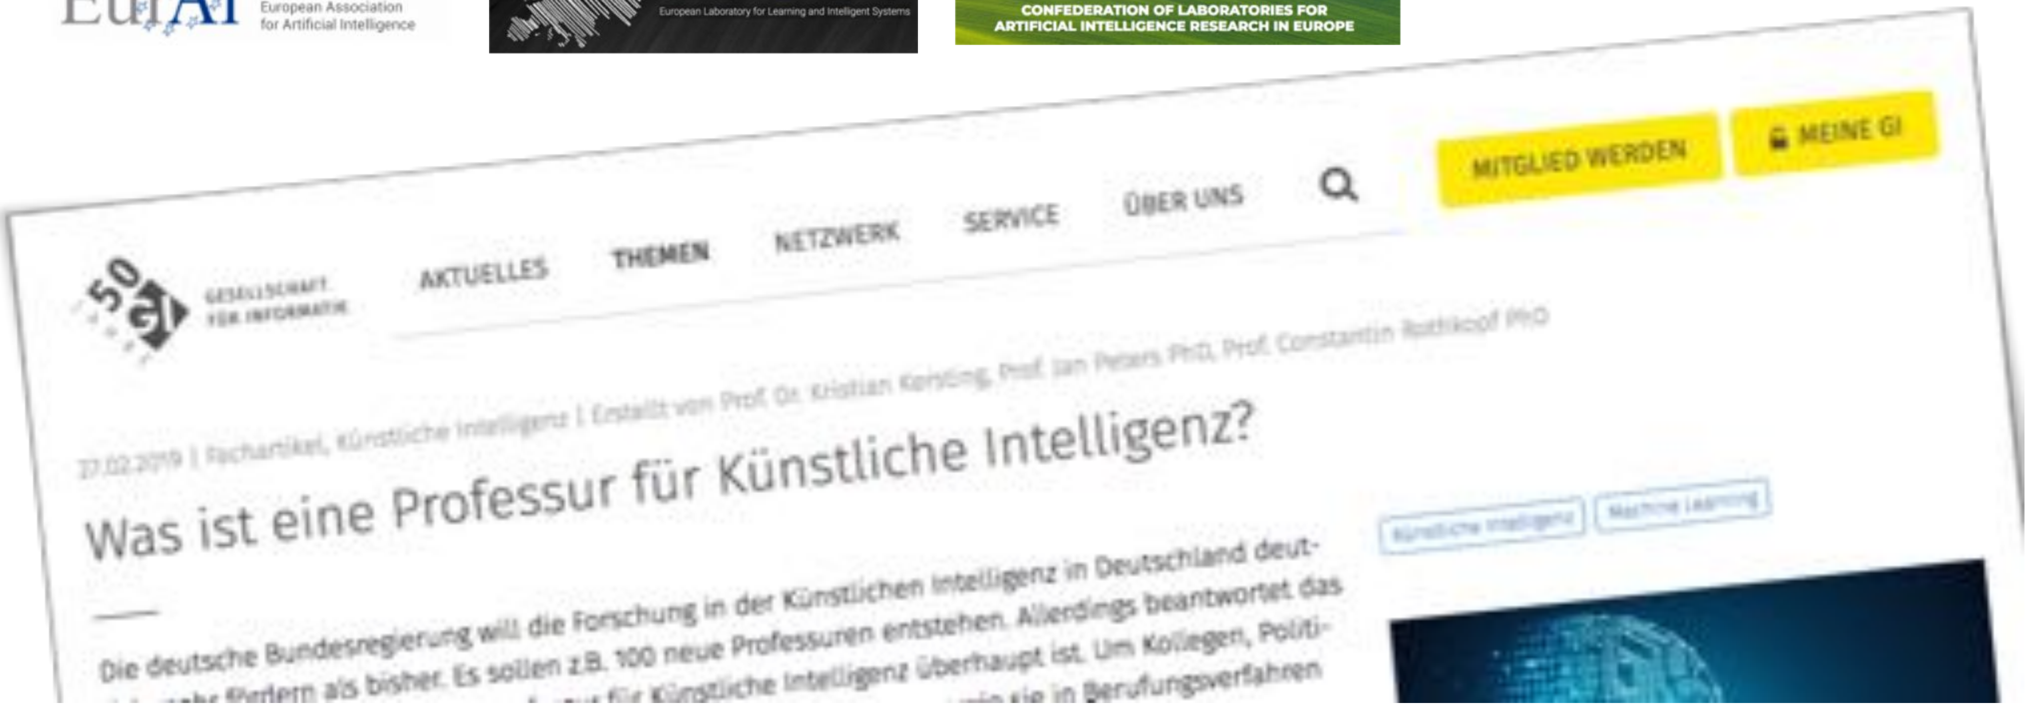

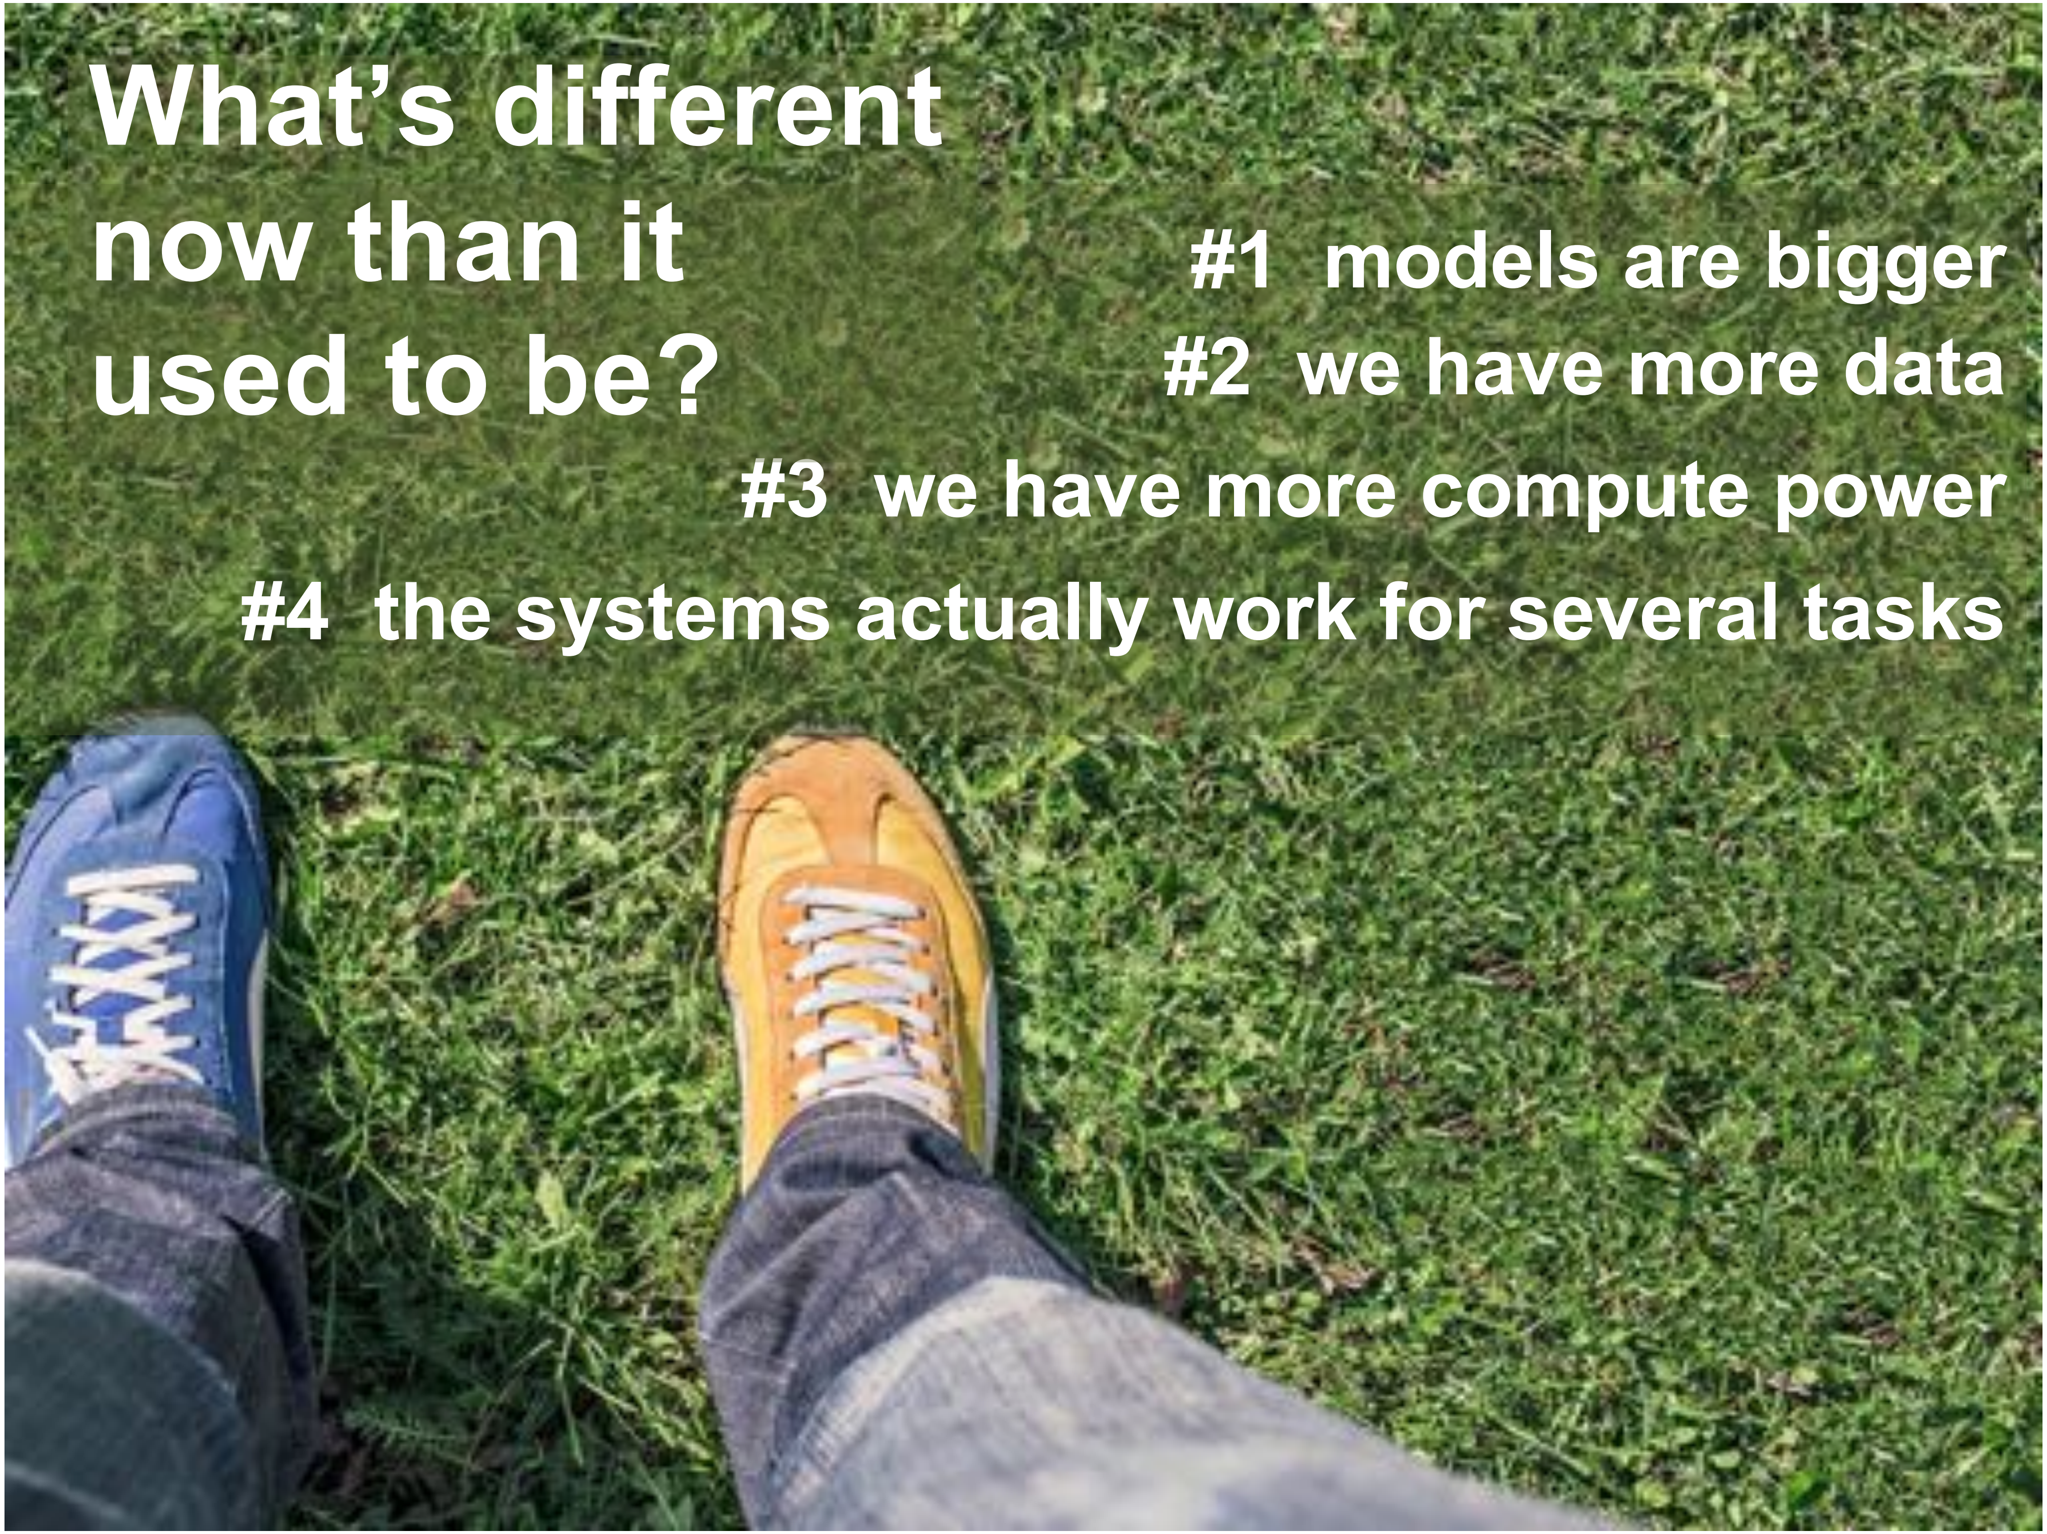
What's different now than it used to be?
#1 models are bigger
#2 we have more data
#3 we have more compute power
#4 the systems actually work for several tasks

AI can learn to manipulate objects

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

Deep Neural Networks

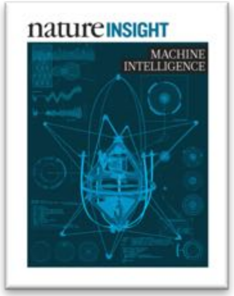

Potentially much more powerful than shallow architectures, represent computations

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

Neuron

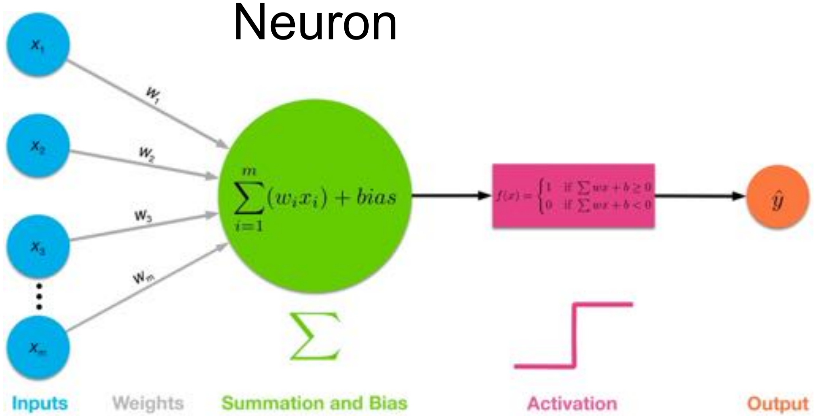

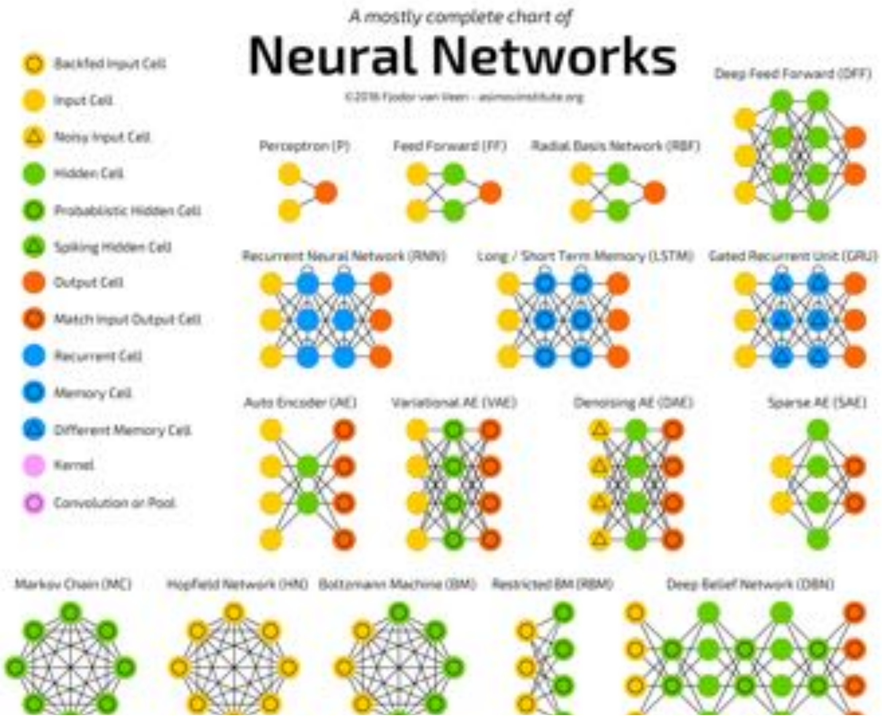

Differentiable Programming

Deep Neural Networks

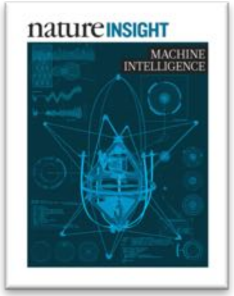

Potentially much more powerful than shallow architectures, represent computations

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

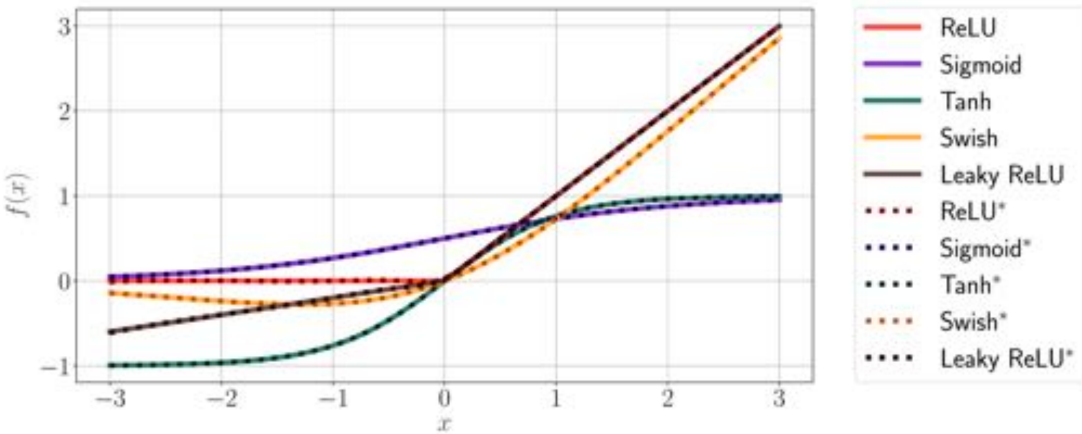

Fashion MNIST

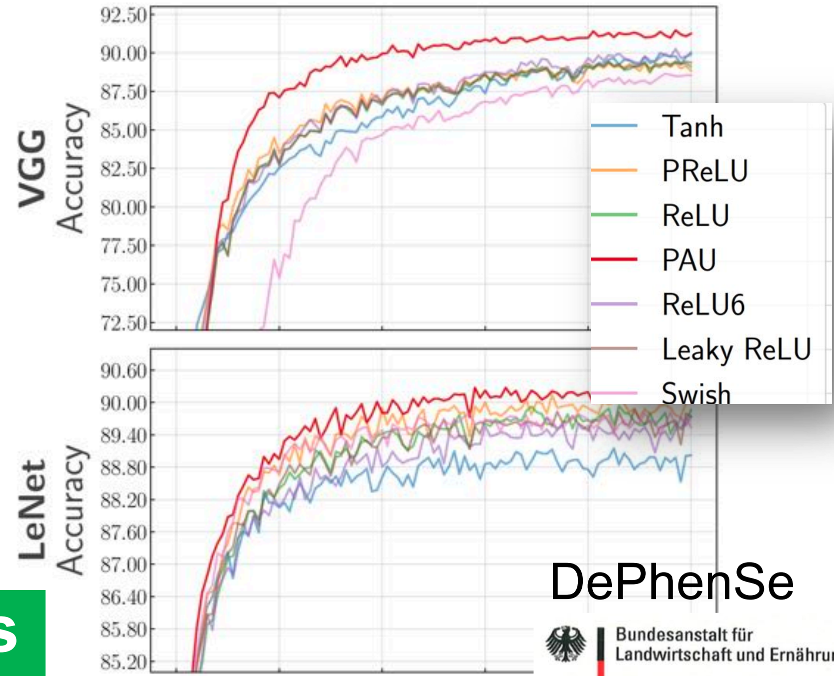

https://github.com/ml-research/pau

E2E-Learning Activation Functions

[Molina, Schramowski, Kersting arxiv:1901.03704 2019]

DePhenSe

Bundesanstalt für Landwirtschaft und Ernährung

Deep Neural Networks

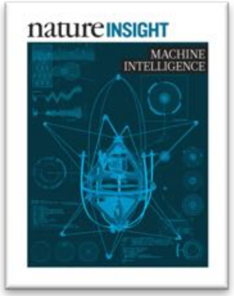

Potentially much more powerful than shallow architectures, represent computations

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

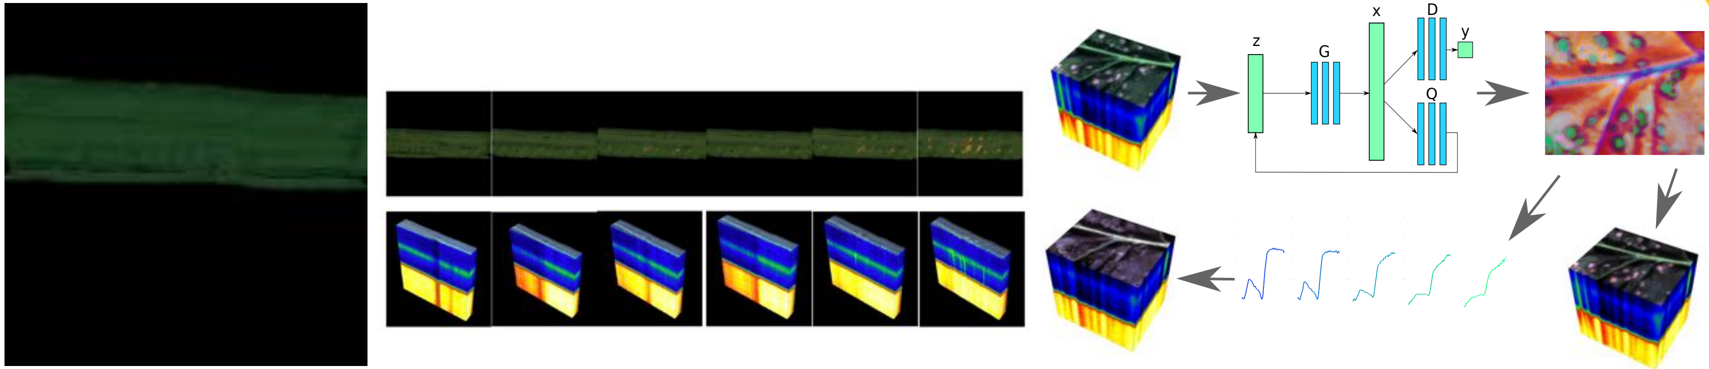

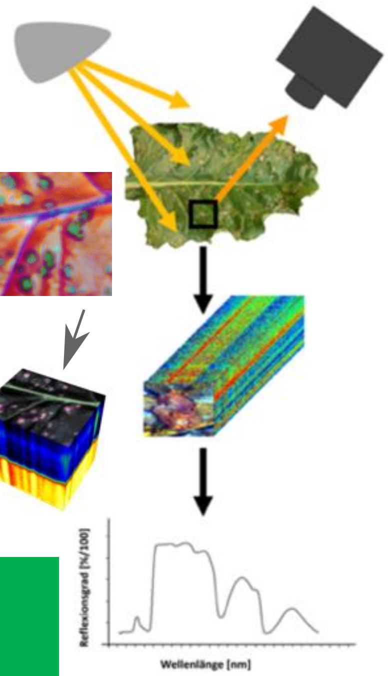

They "develop intuition" about complicated biological processes and generate scientific data

[Schramowski, Brugger, Mahlein, Kersting 2019]

DePhenSe

Bundesanstalt für Landwirtschaft und Ernährung

Deep Neural Networks

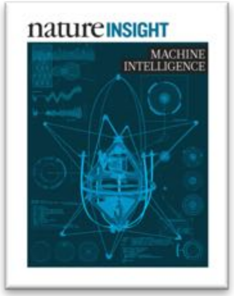

Potentially much more powerful than shallow architectures, represent computations

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

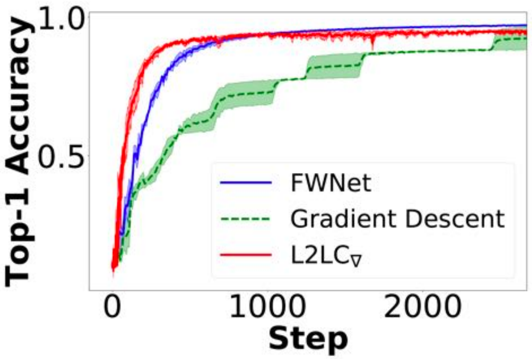

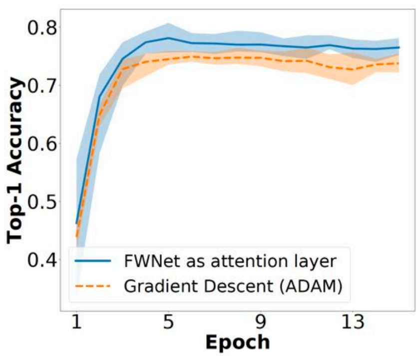

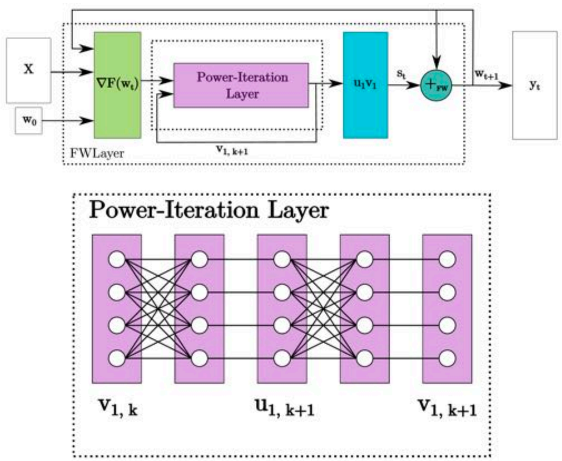

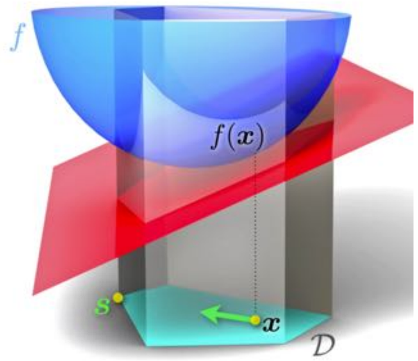

They "invent" constrained optimizers

DePhenSe

Bundesanstalt für Landwirtschaft und Ernährung

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

Deep Neural Networks

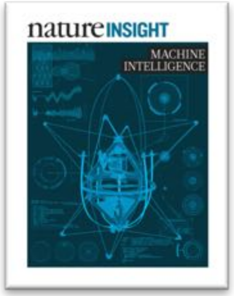

Potentially much more powerful than shallow architectures, represent computations

[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

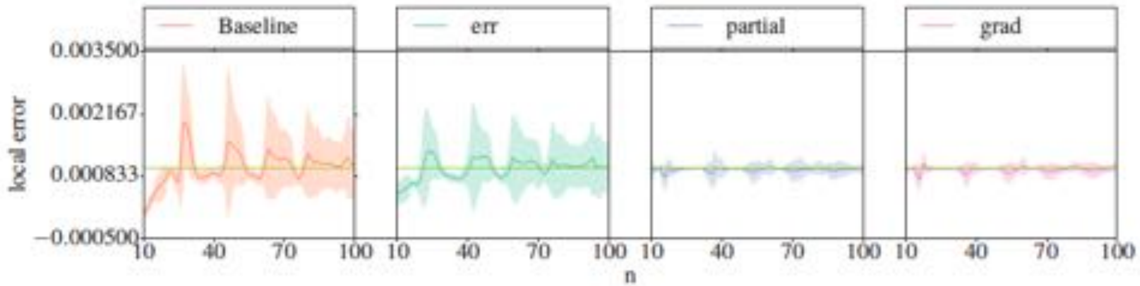

They can learn to integrate

[Jentzsch, Schramowski, Kersting to be submitted 2019]

DePhenSe

Bundesanstalt für Landwirtschaft und Ernährung

Deep Neural Networks

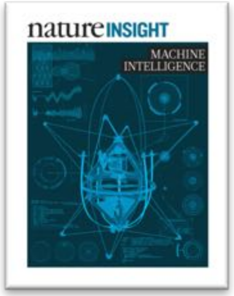

Potentially much more powerful than shallow architectures, represent computations

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

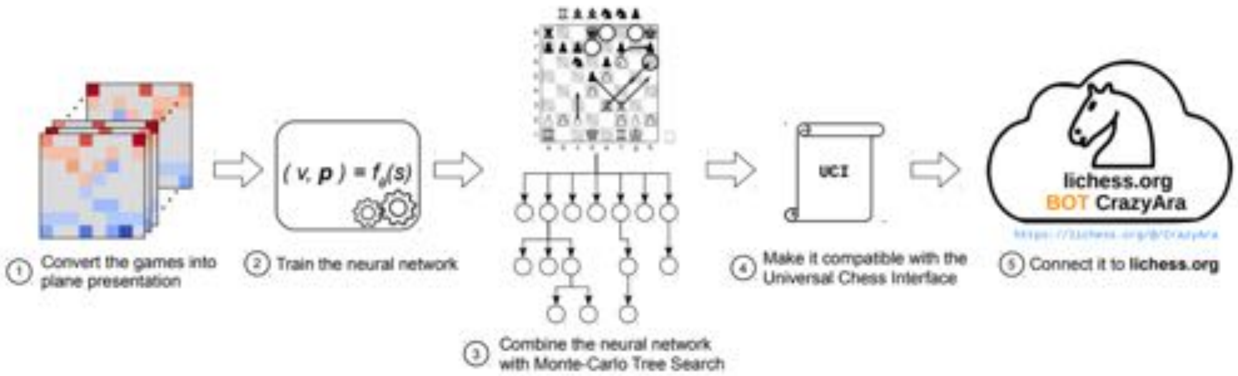

They can beat the world champion in CrazyHouse

[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)

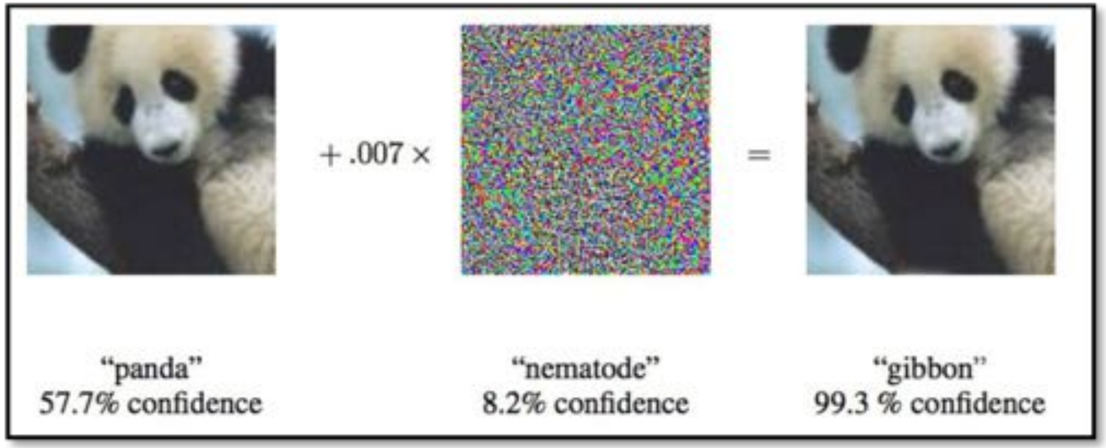

Google, 2015

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

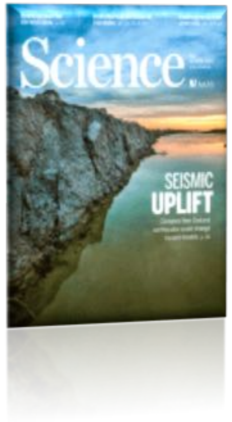

The Quest for a „good“ AI

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?

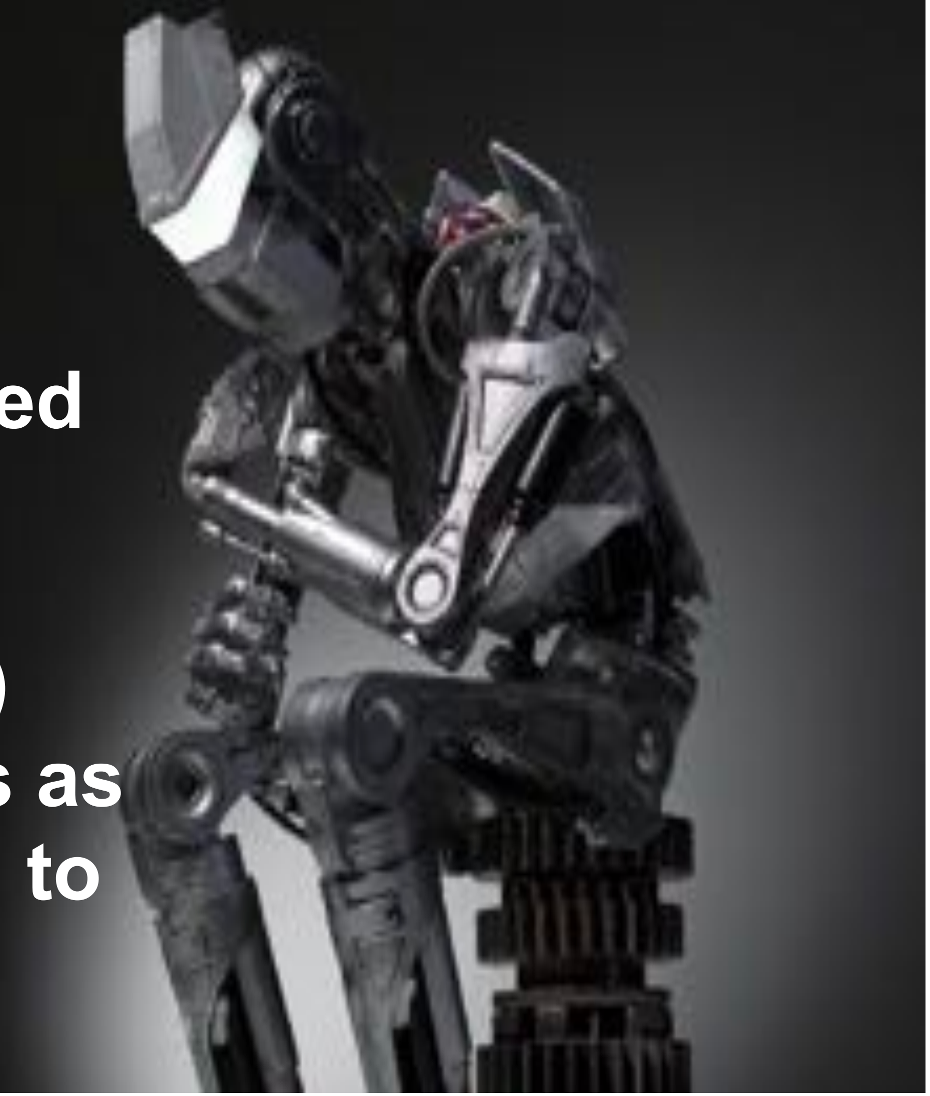

„The Ethics of Artificial Intelligence“ Cambridge Handbook of Artificial Intelligence, 2011

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

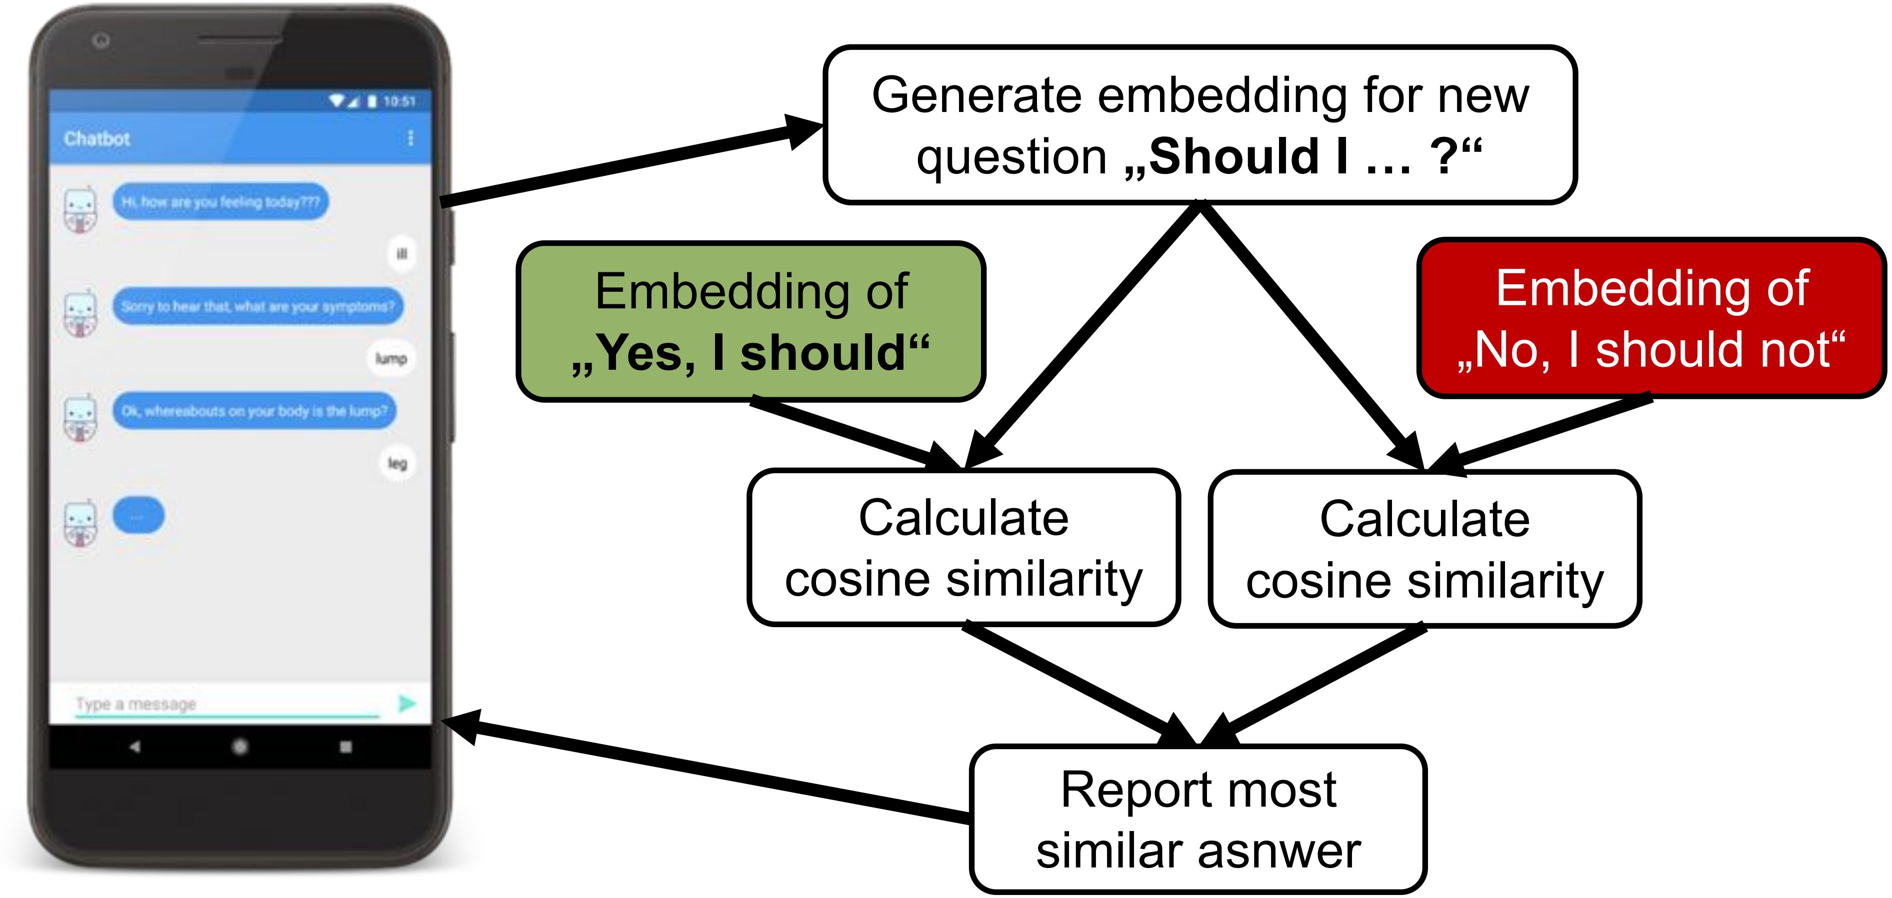

The Moral Choice Machine

Not all stereotypes are bad

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

Can we trust deep neural networks?

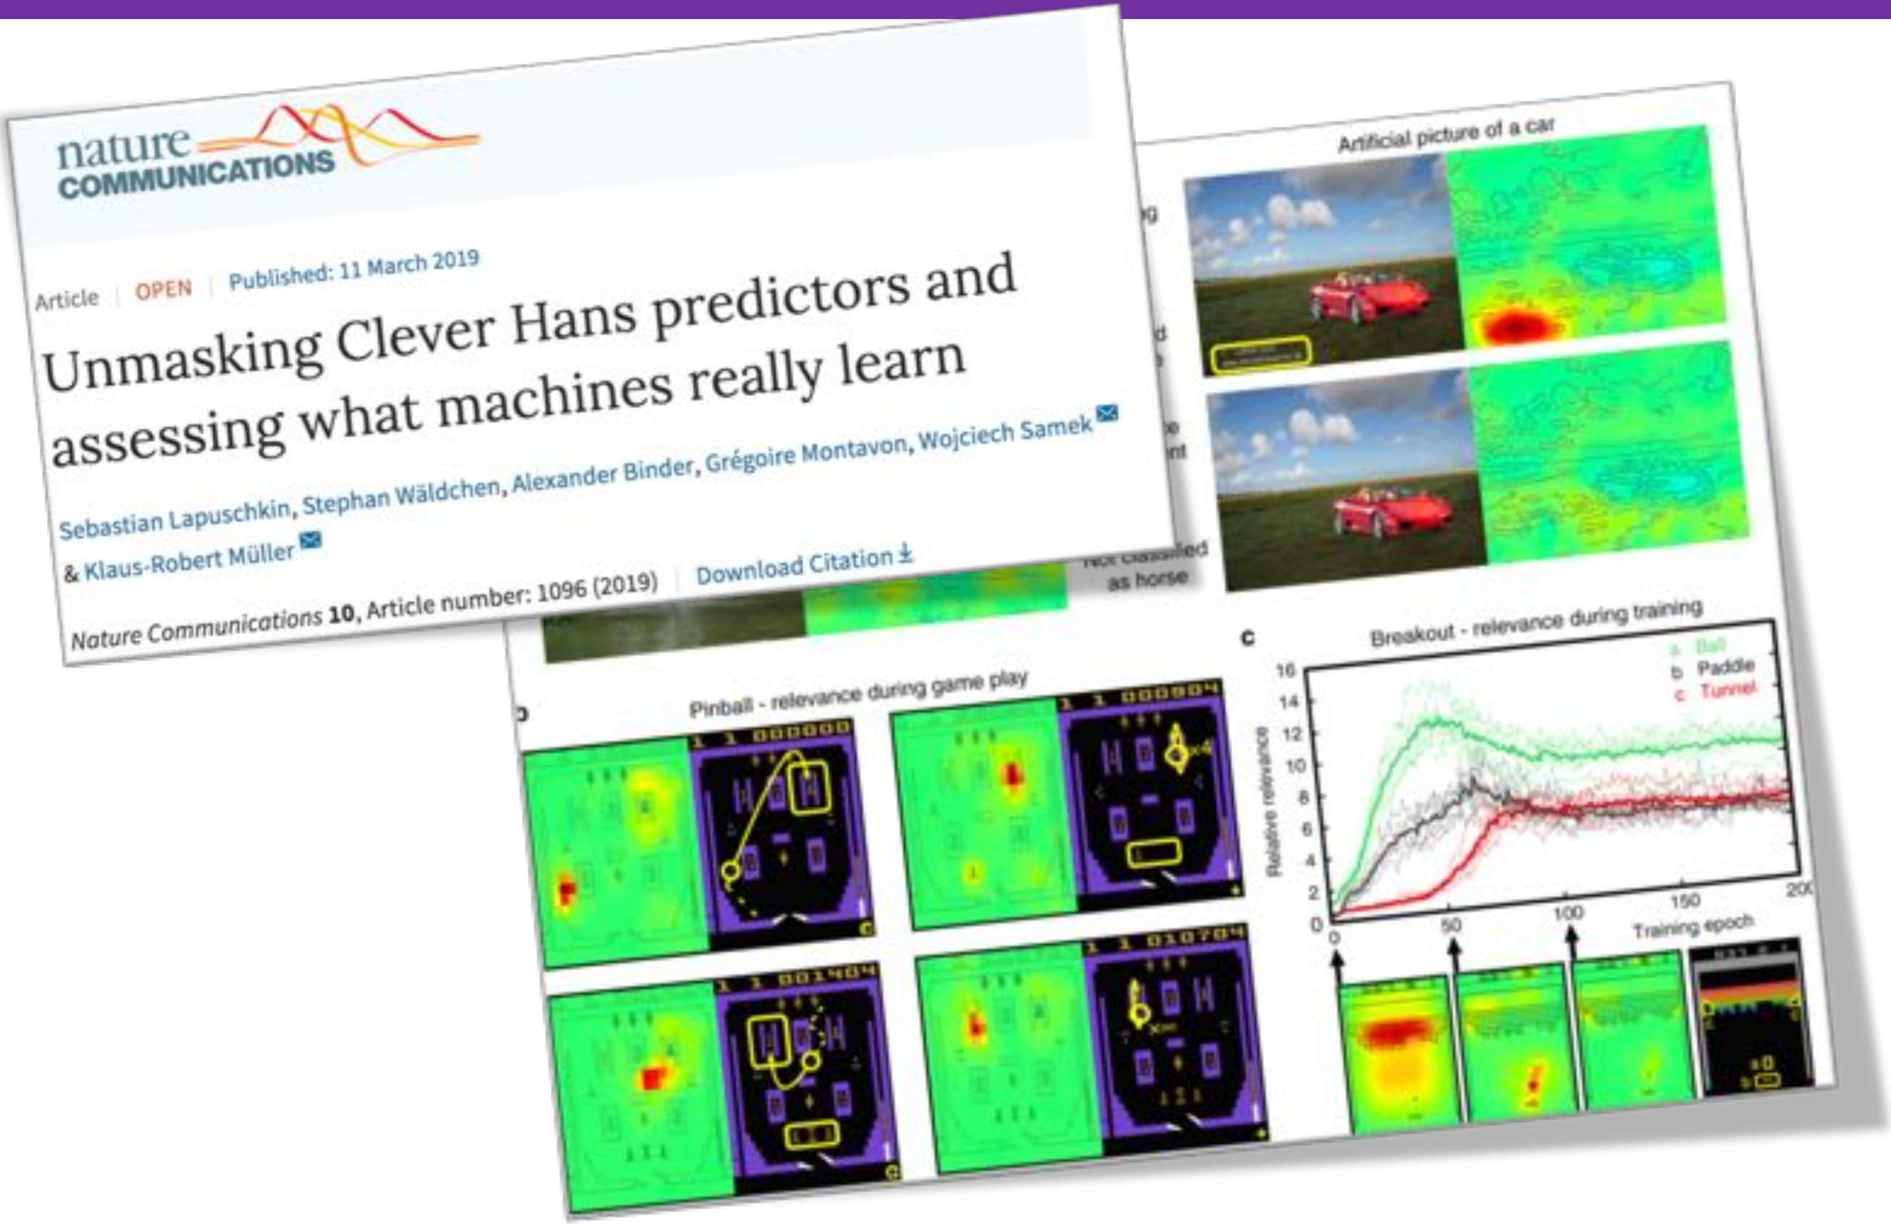

DNNs often have no probabilistic semantics. They are not calibrated joint distributions.

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

MNIST

Train & Evaluate

SVHN

SEMEION

Transfer Testing

[Bradshaw et al. arXiv:1707.02476 2017]

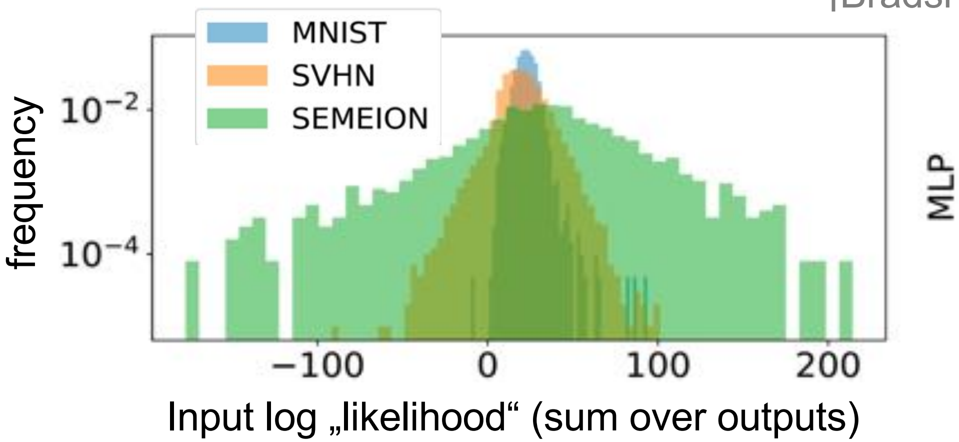

Many DNNs cannot distinguish the datasets

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

The third wave of differentiable programming

Getting deep systems that know when they do not know and, hence, recognise new situations and adapt to them

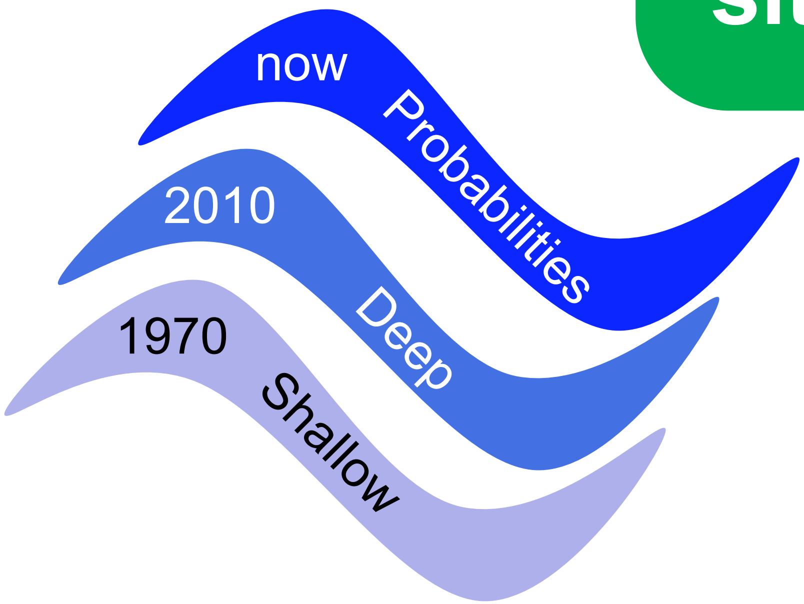

This results in Sum-Product Networks, a deep probabilistic learning framework

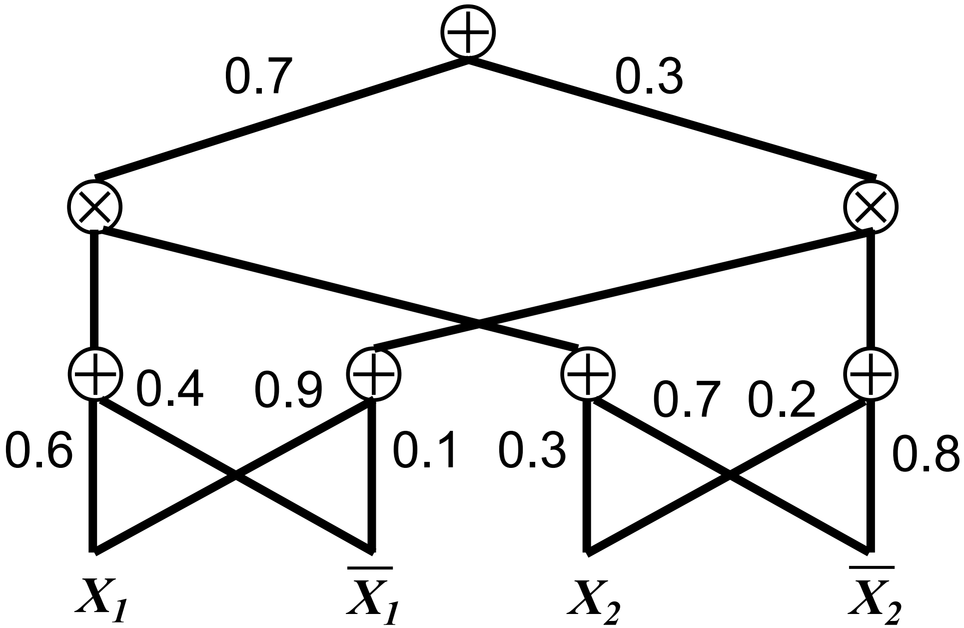

Computational graph (kind of TensorFlow graphs) that encodes how to compute probabilities

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]

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
Max Planck Institute for Intelligent Systems
UNIVERSITY OF CAMBRIDGE
VECTOR INSTITUTE
CAML
MADESI
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

```
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

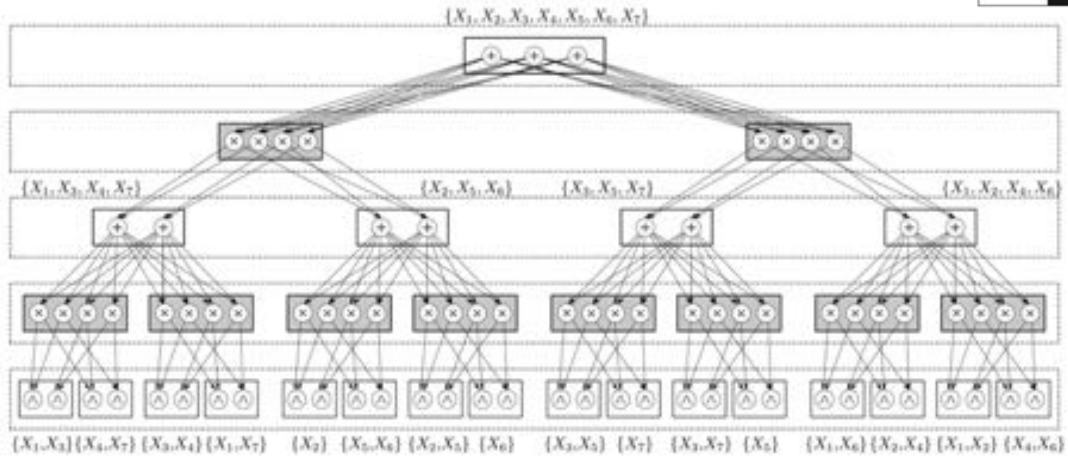

Build a random SPN structure. This can be done in an informed way or completely at random

		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)

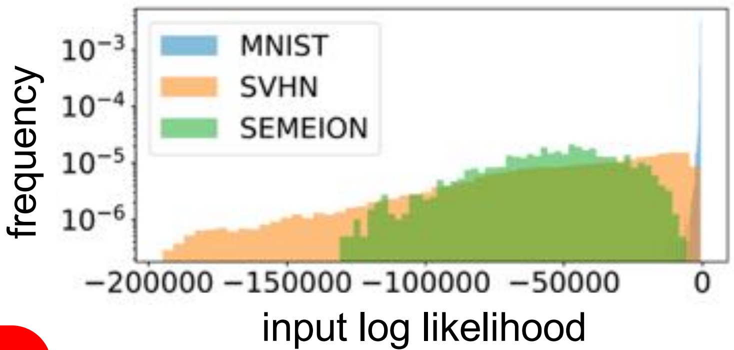

outliers
prototypes
outliers
prototypes

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]

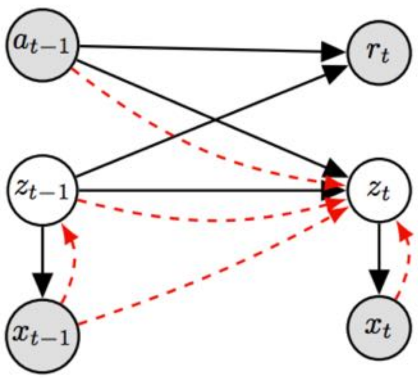

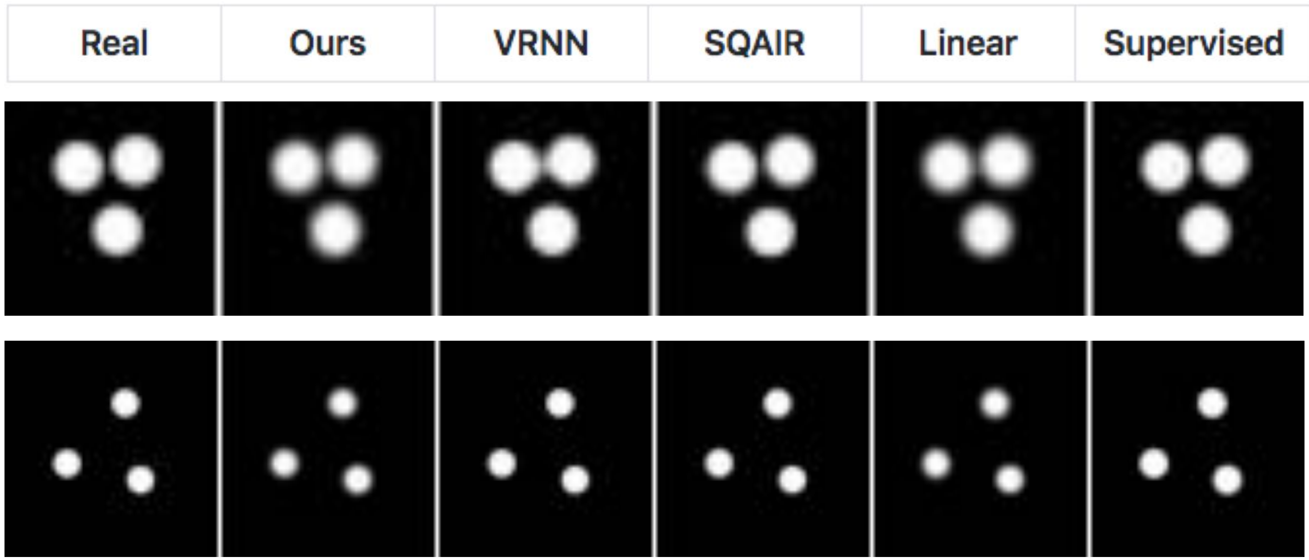

putting structure and tractable inference into deep models

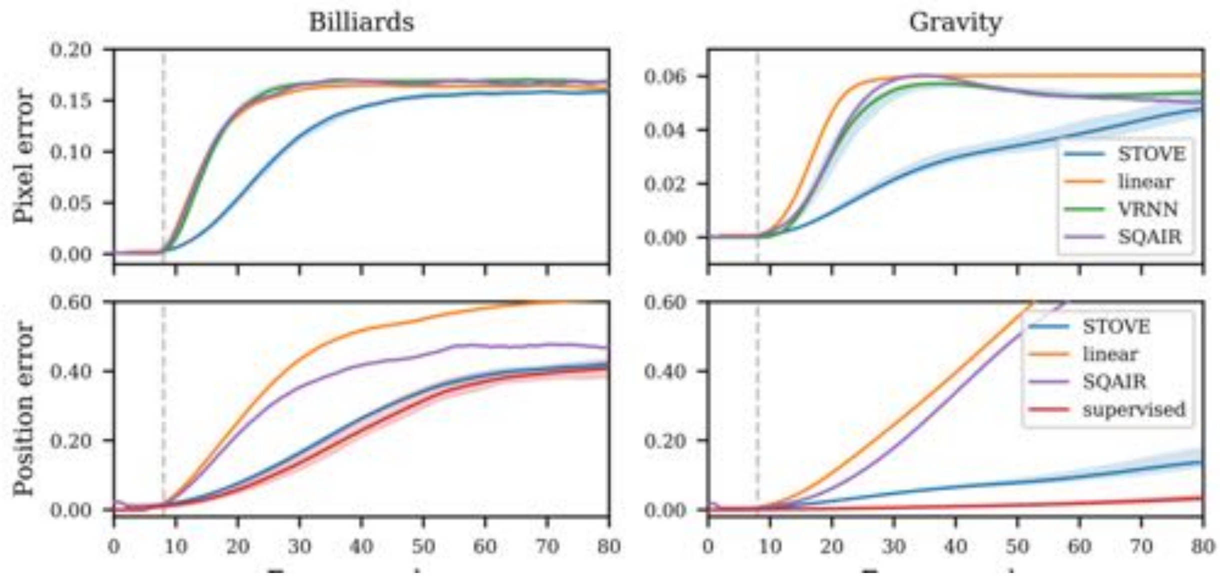

However, there are not enough data scientists, statisticians, machine learning and AI experts

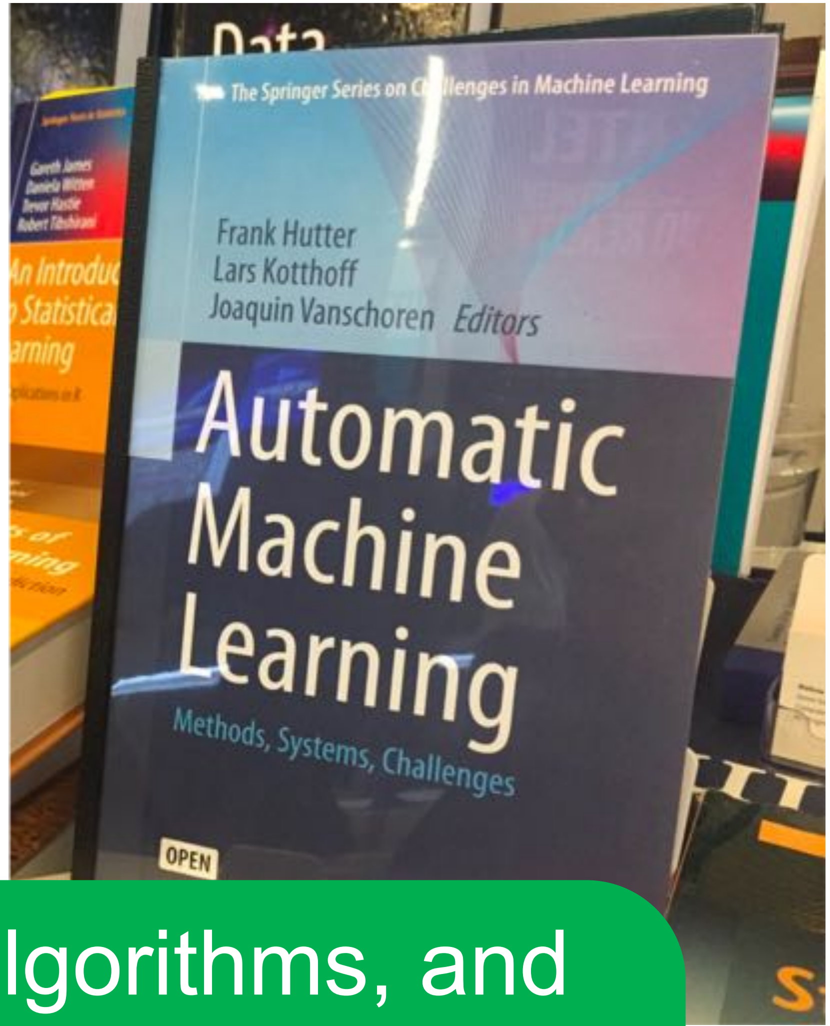

Provide the foundations, algorithms, and tools to develop systems that ease and support building ML/AI models as much as possible and in turn help reproducing and hopfeully even justifying our results

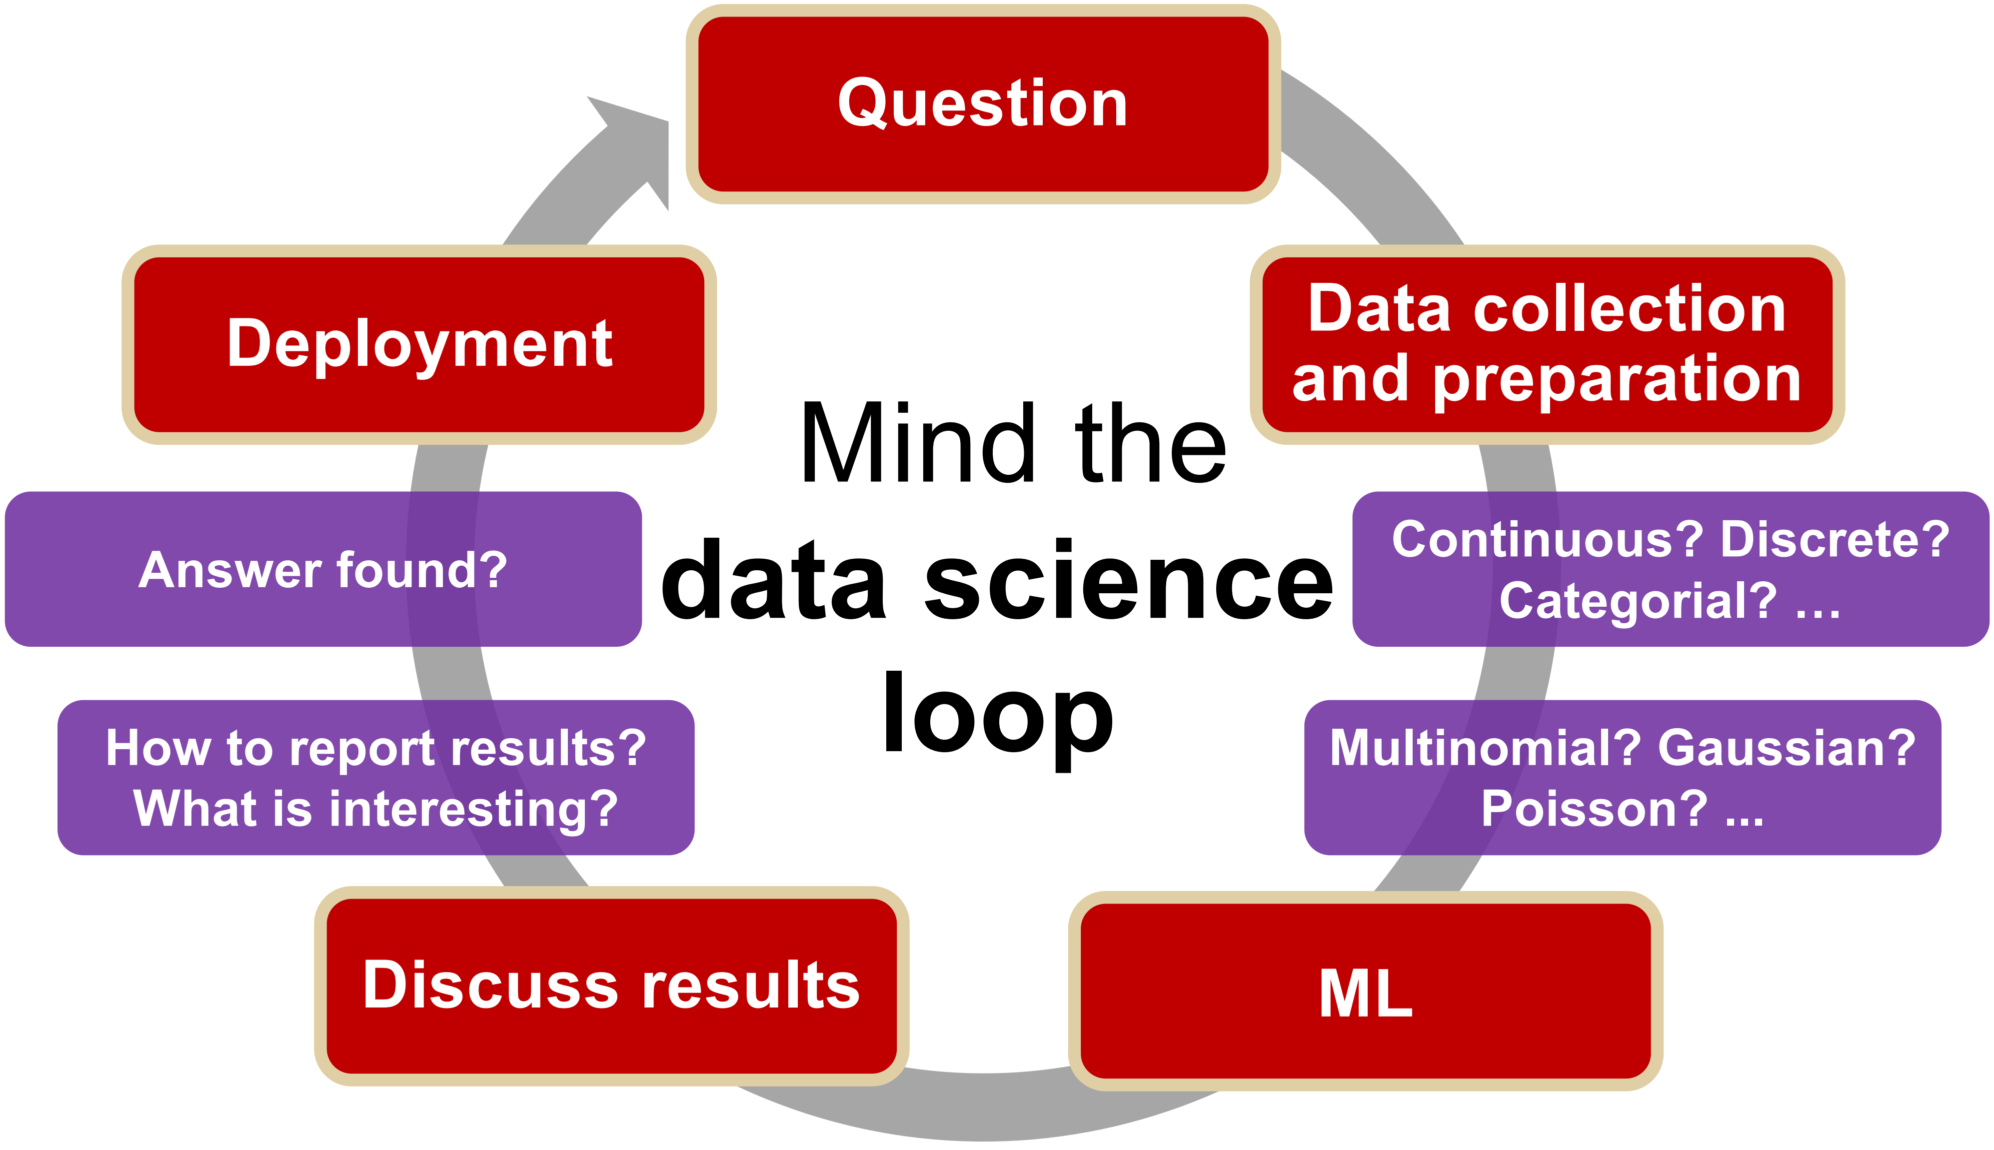
Question
Data collection and preparation
Continuous? Discrete? Categorial? …
Multinomial? Gaussian? Poisson? …
ML
Discuss results
How to report results? What is interesting?
Answer found?
Deployment
Mind the data science loop

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

The Automatic Data Scientist

missing value

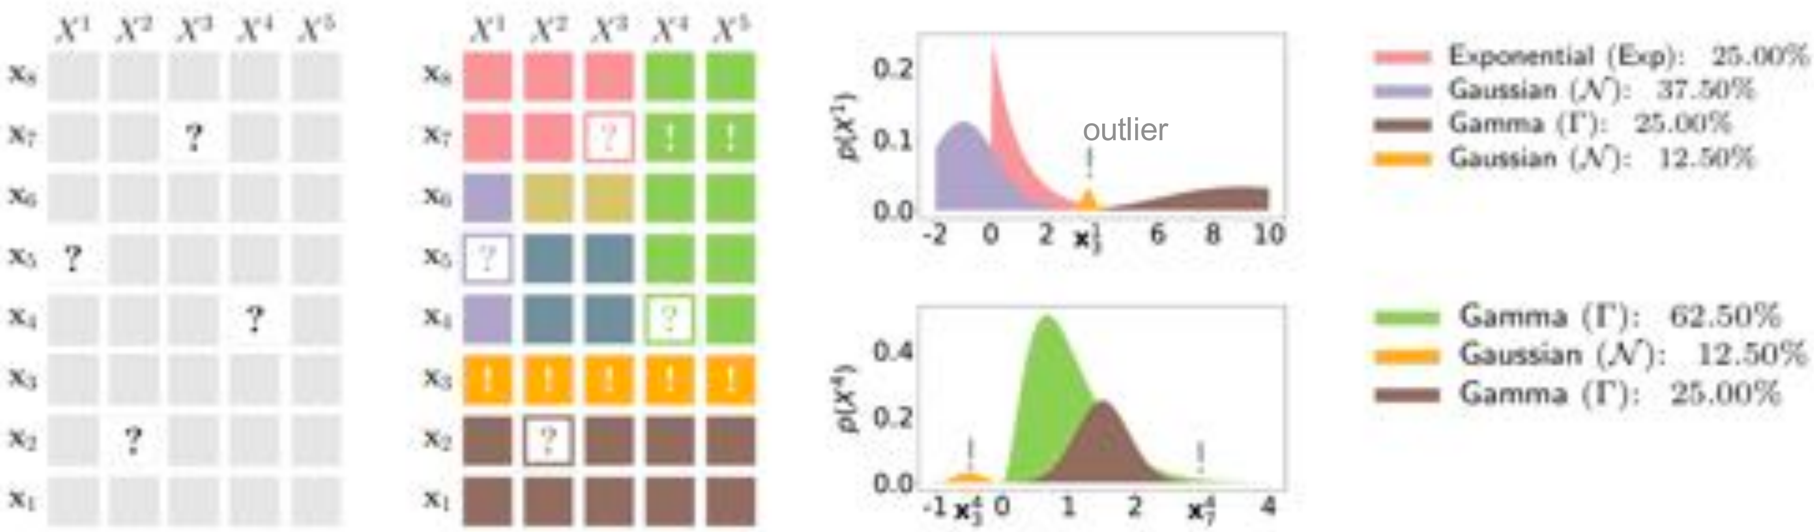

We can even automatically discovers the statistical types and parametric forms of the variables

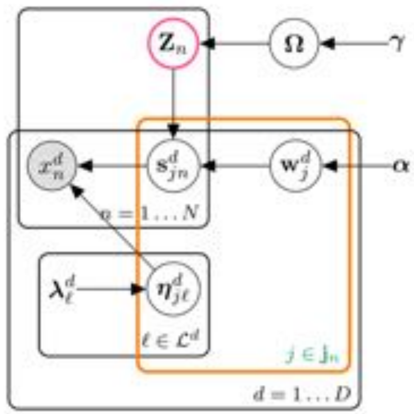

Bayesian Type Discovery

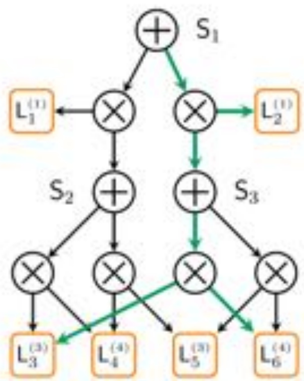

Mixed Sum-Product Network

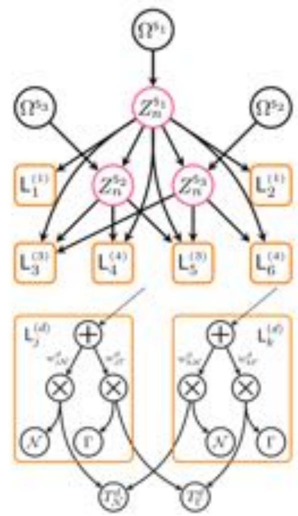

Automatic Statistician

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

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

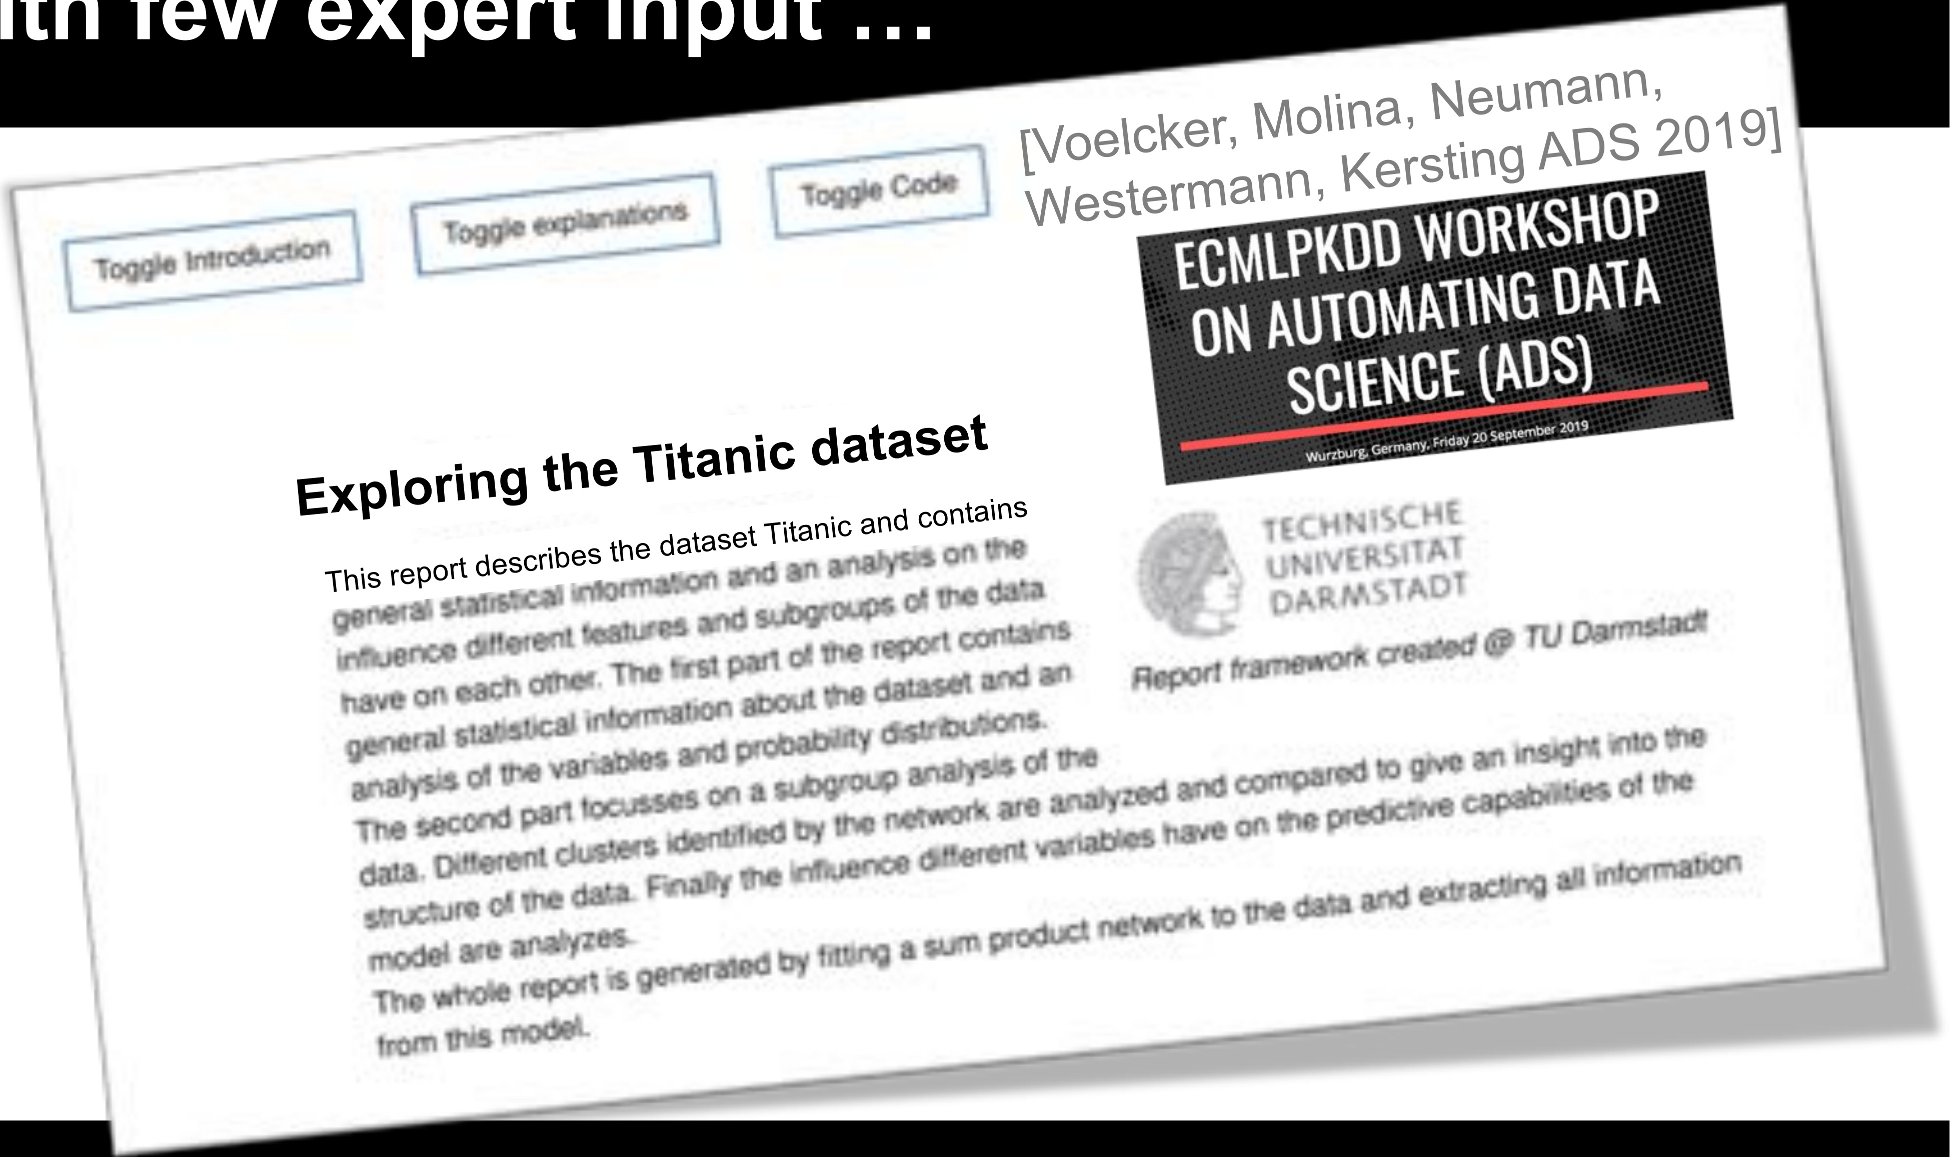

…and can compile data reports automatically

UNIVERSITY OF CAMBRIDGE

UBER AI Labs

Max Planck Institute for Intelligent Systems

UNI GRAZ

FLOW

TECHNISCHE UNIVERSITÄT DARMSTADT

Microsoft Research

UBER AI Labs

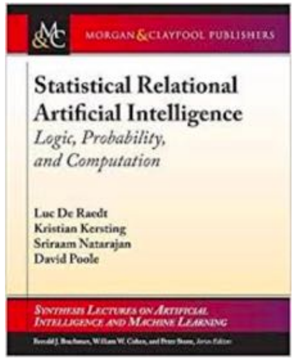

relationalAI

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"

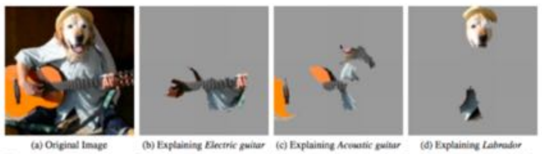

Figure 4: Explaining an image classification prediction made by Google's Inception network, highlighting positive pixels. The top 3 classes predicted are "Electric Guitar" ($p = 0.32$), "Acoustic guitar" ($p = 0.24$) and "Labrador" ($p = 0.21$)

Teso, Kersting AIES 2019

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

Making Clever Hans Clever

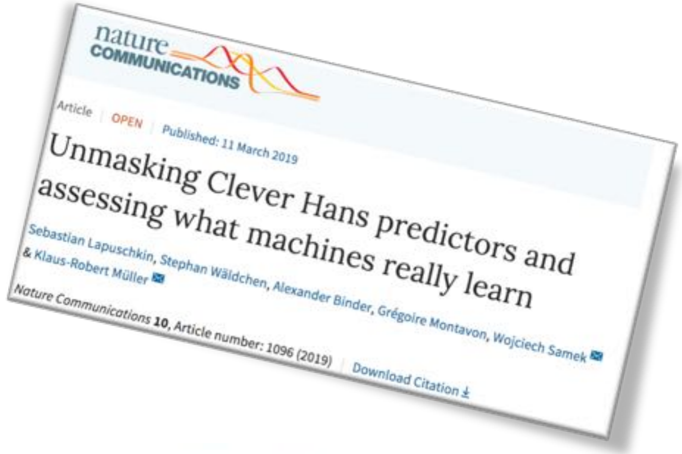
nature COMMUNICATIONS

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

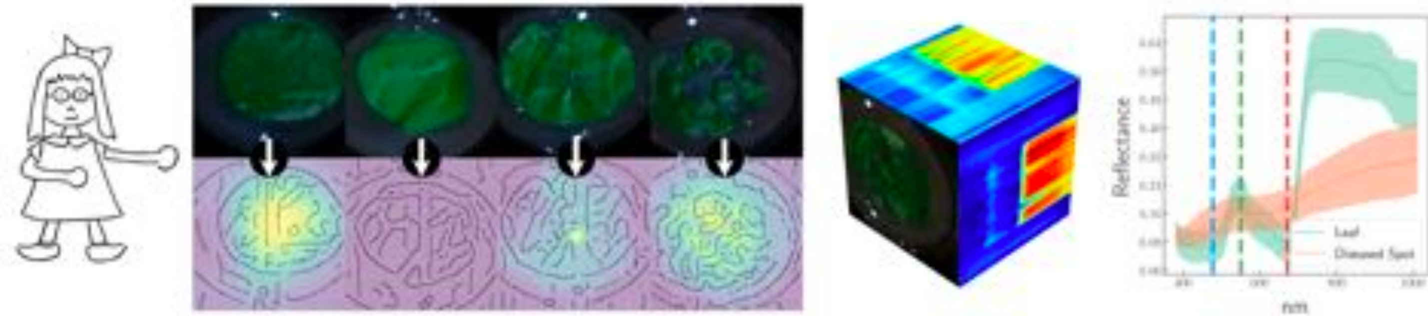

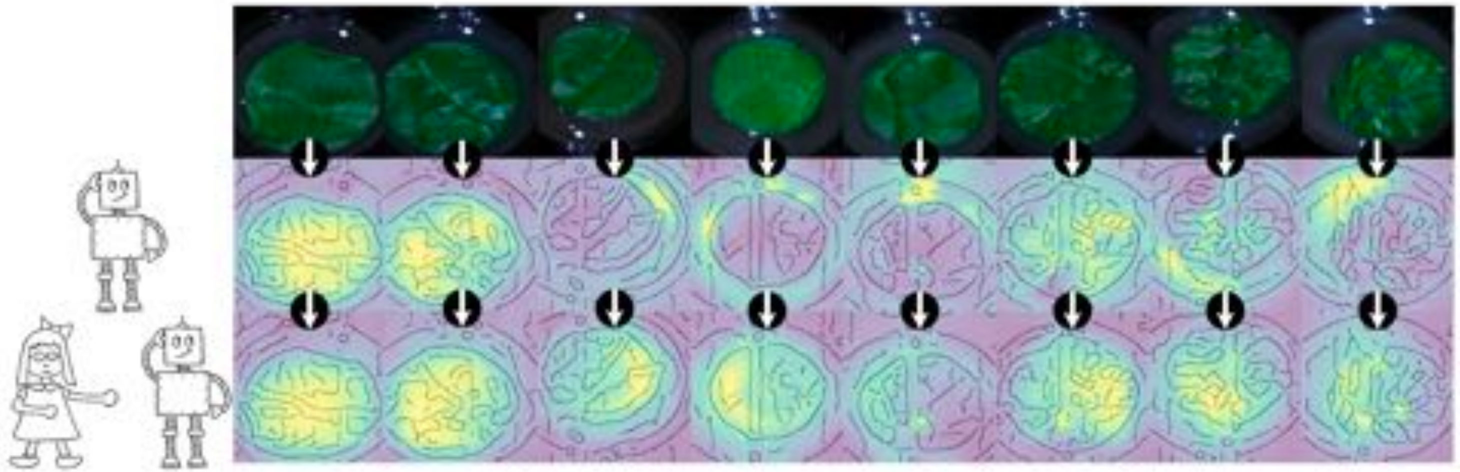

[Teso, Kersting AIES 2019 and ongoing research]

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

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

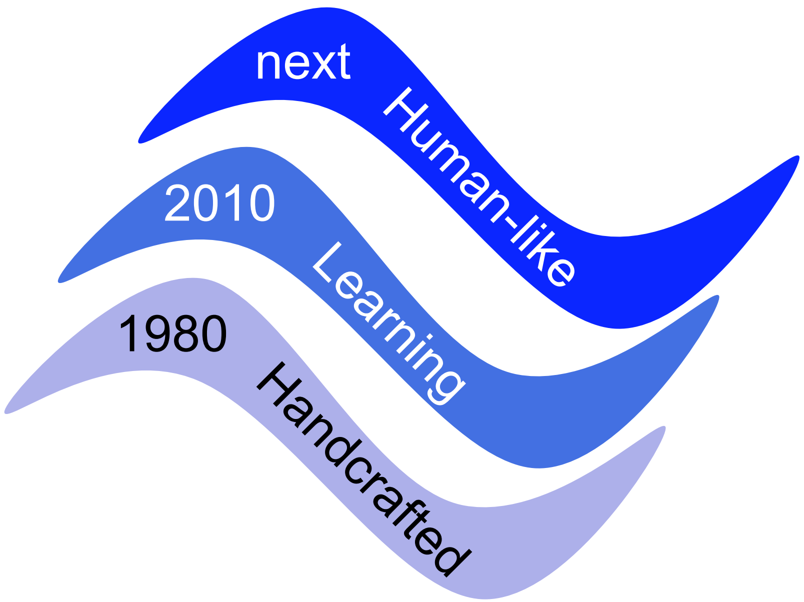

AI systems that can acquire human-like communication and reasoning capabilities, with the ability to recognise new situations and adapt to them.

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