The Third Wave of AI

TECHNISCHE UNIVERSITÄT DARMSTADT

Prof. Dr. Kristian Kersting

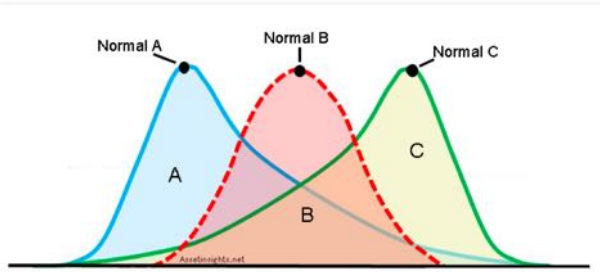

Getting deep systems that reason and know when they don't know

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

My team and I in the Machine Learning group would like to make computers learn so much about the world, so rapidly and flexibly, as humans.

TECHNISCHE UNIVERSITÄT DARMSTADT

Lernende Systeme
DIE PLATTFORM FÜR KÜNSTLICHE INTELLIGENZ

Federal Ministry of Education and Research

2017 - now:	Professor (W3) for Machine Learning at the CS Department of the TU Darmstadt, Germany
2013 - 2017:	Associate Professor (W2) for Data Mining at the CS Department of the TU Dortmund University, Germany
2012 - 2013:	Assistant Professor (W1) for Spatio-Temporal Pattern in Agriculture at the Faculty of Agriculture of the University of Bonn, Germany
2008 - 2012:	Fraunhofer Attract research group leader at the Fraunhofer IAIS, Germany
2007:	PostDoctoral Associate at MIT CSAIL, USA, working with Leslie Kaelbling, Josh Tenenbaum, and Nicholas Roy.
2000 - 2006:	Ph.D. student at the CS Department of the University of Freiburg, Germany, working with Luc De Raedt (supervsior) and Wolfram Burgard.
1996 - 2000:	Diploma in Computer Science at the CS Department of the University of Freiburg, Germany

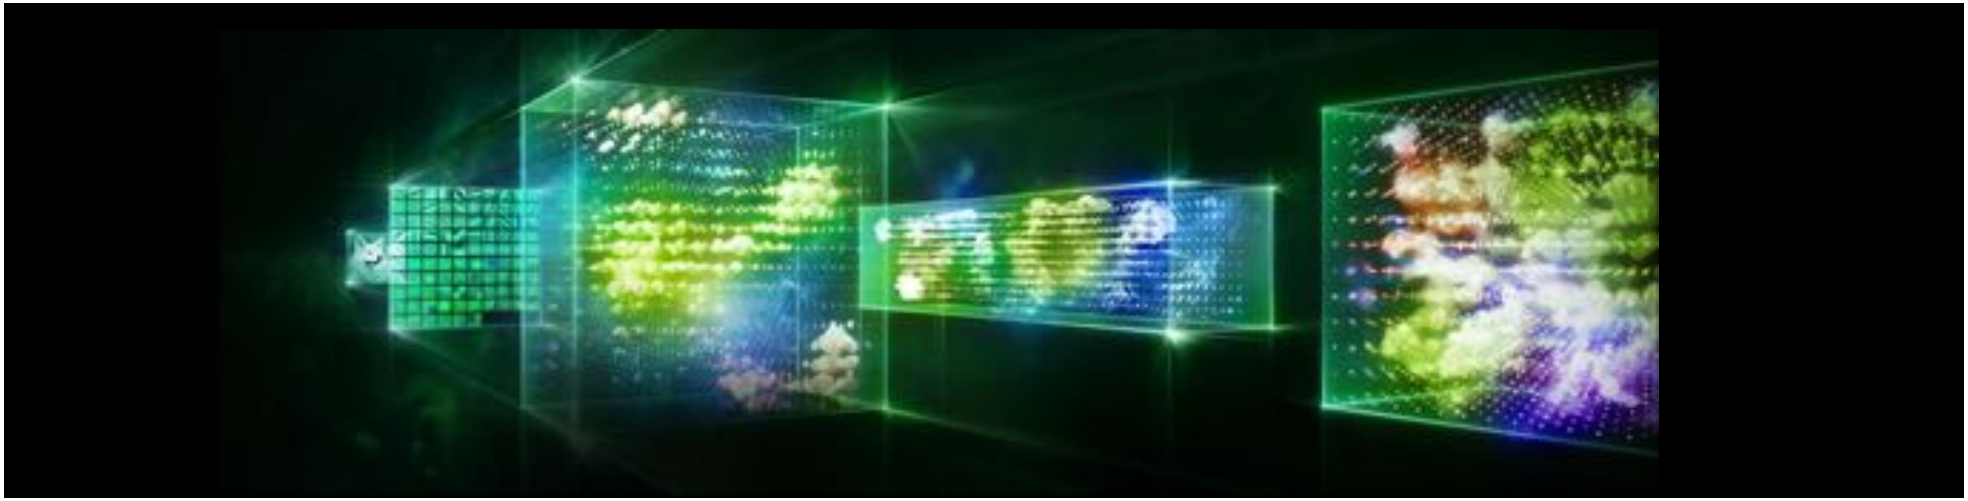

Deep Probabilistic Learning

UBER AI Labs

ABB

The dream of AI is not new

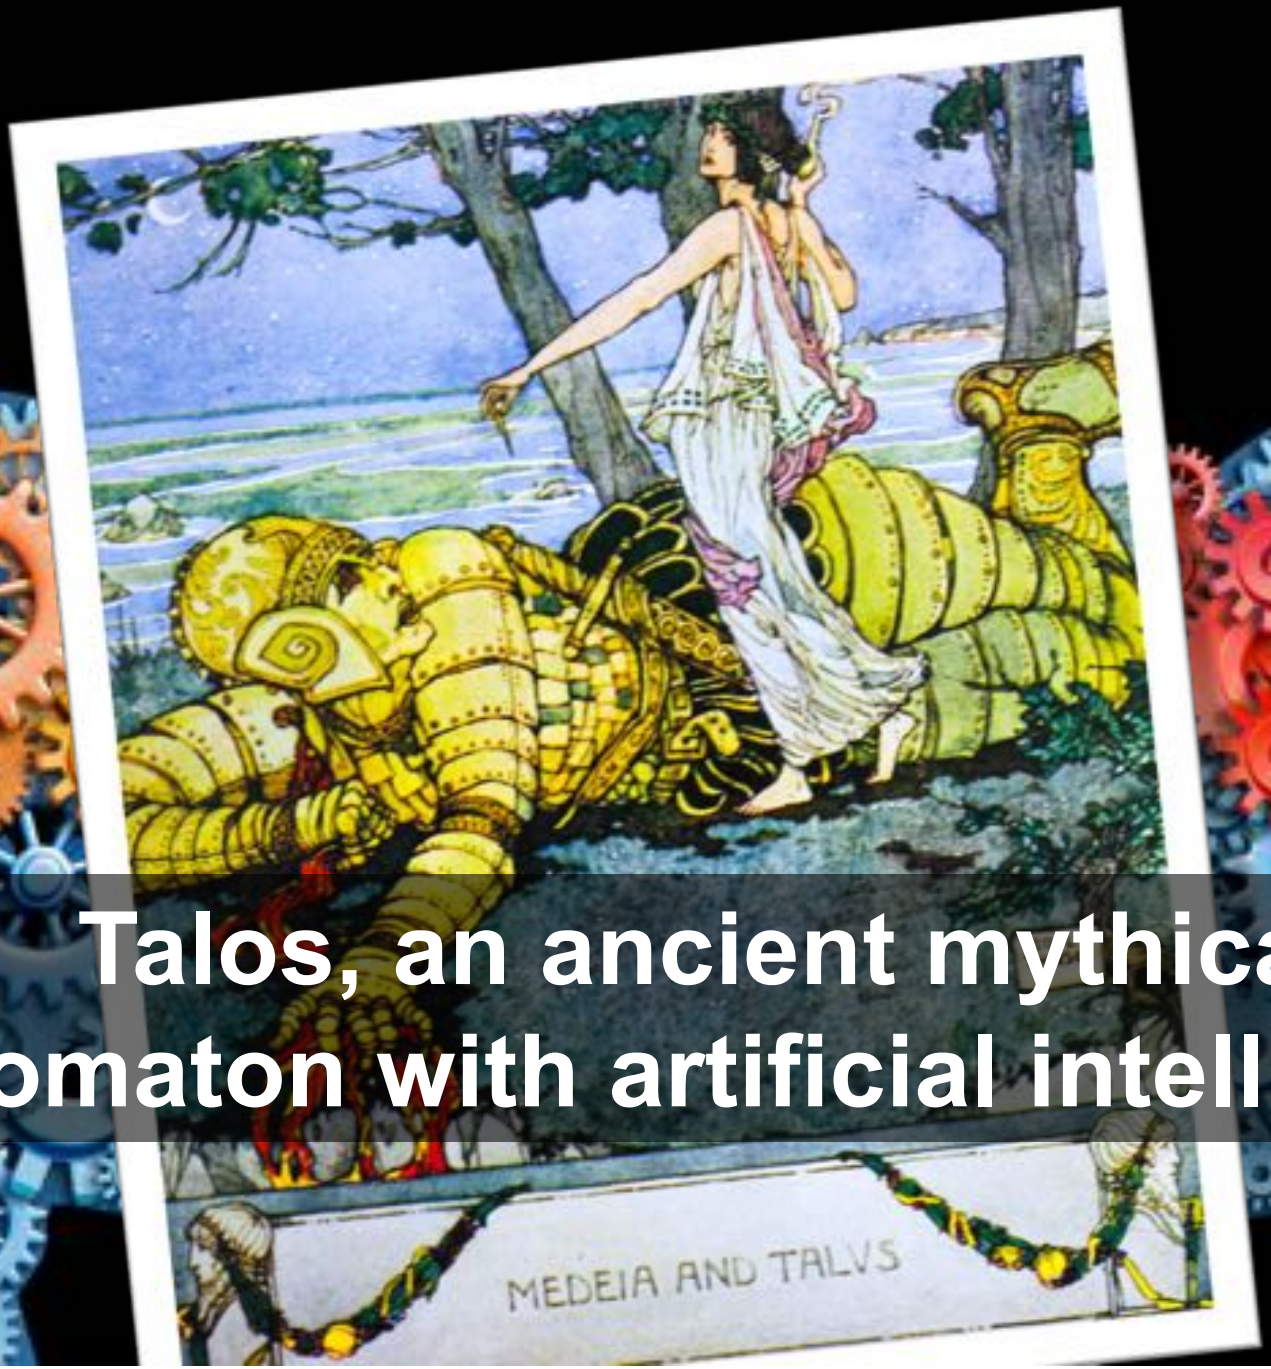

Talos, an ancient mythical automaton with artificial intelligence

AI seems to have many faces

What is AI?

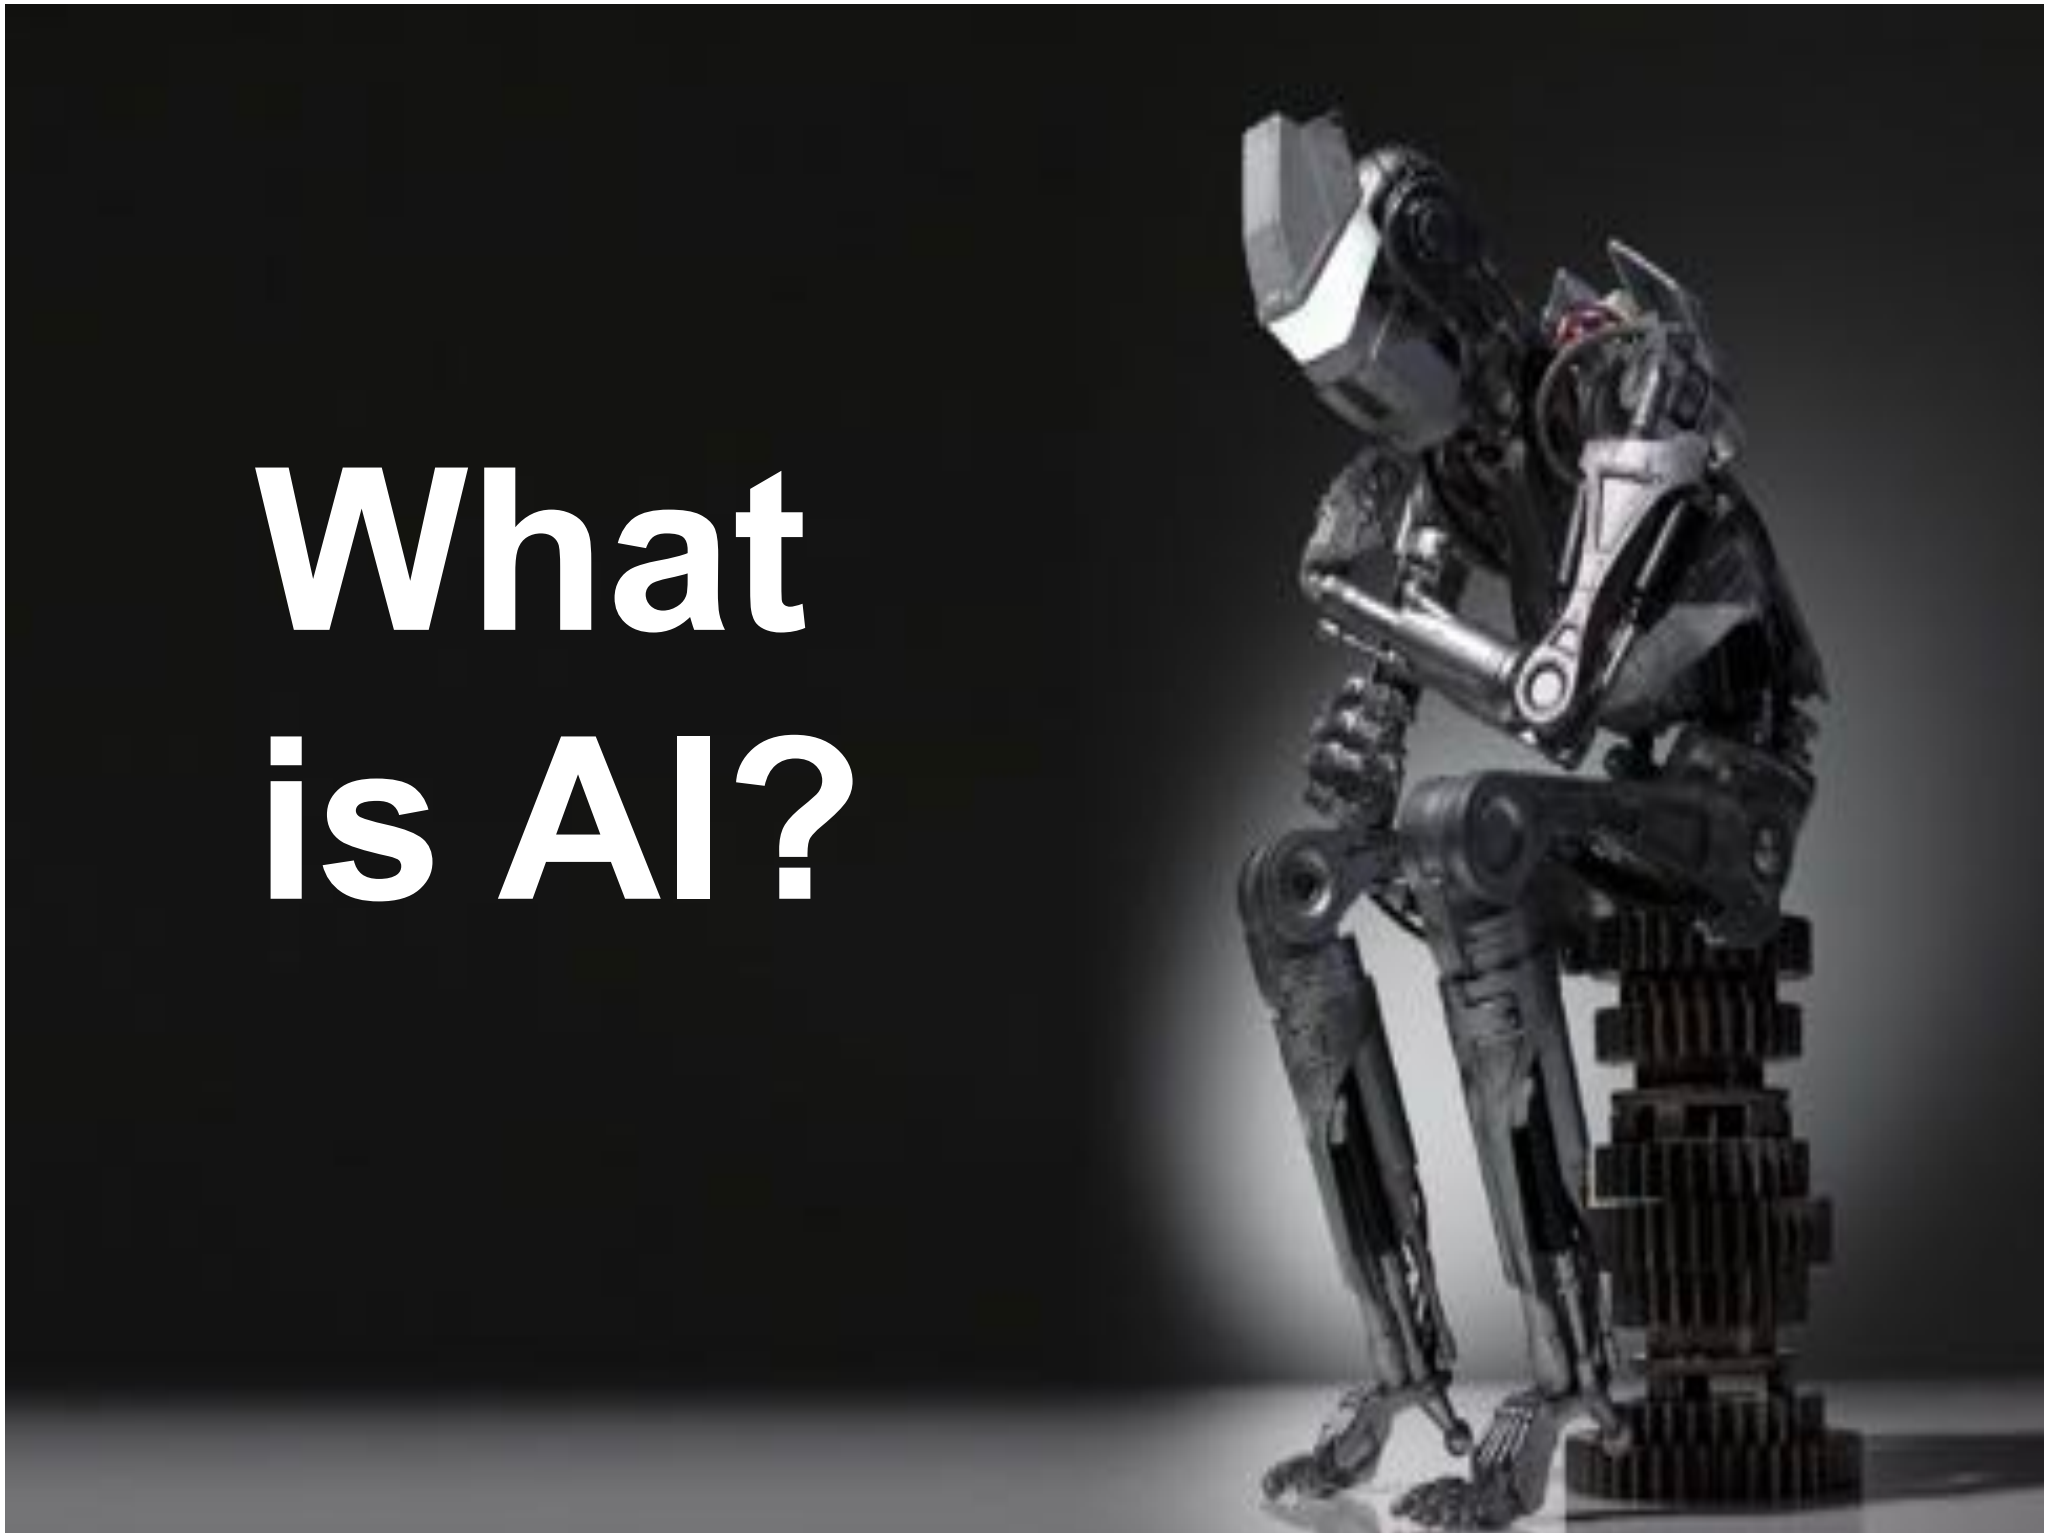

Humans are smart

https://www.youtube.com/watch?v=XQ79UUlOeWc

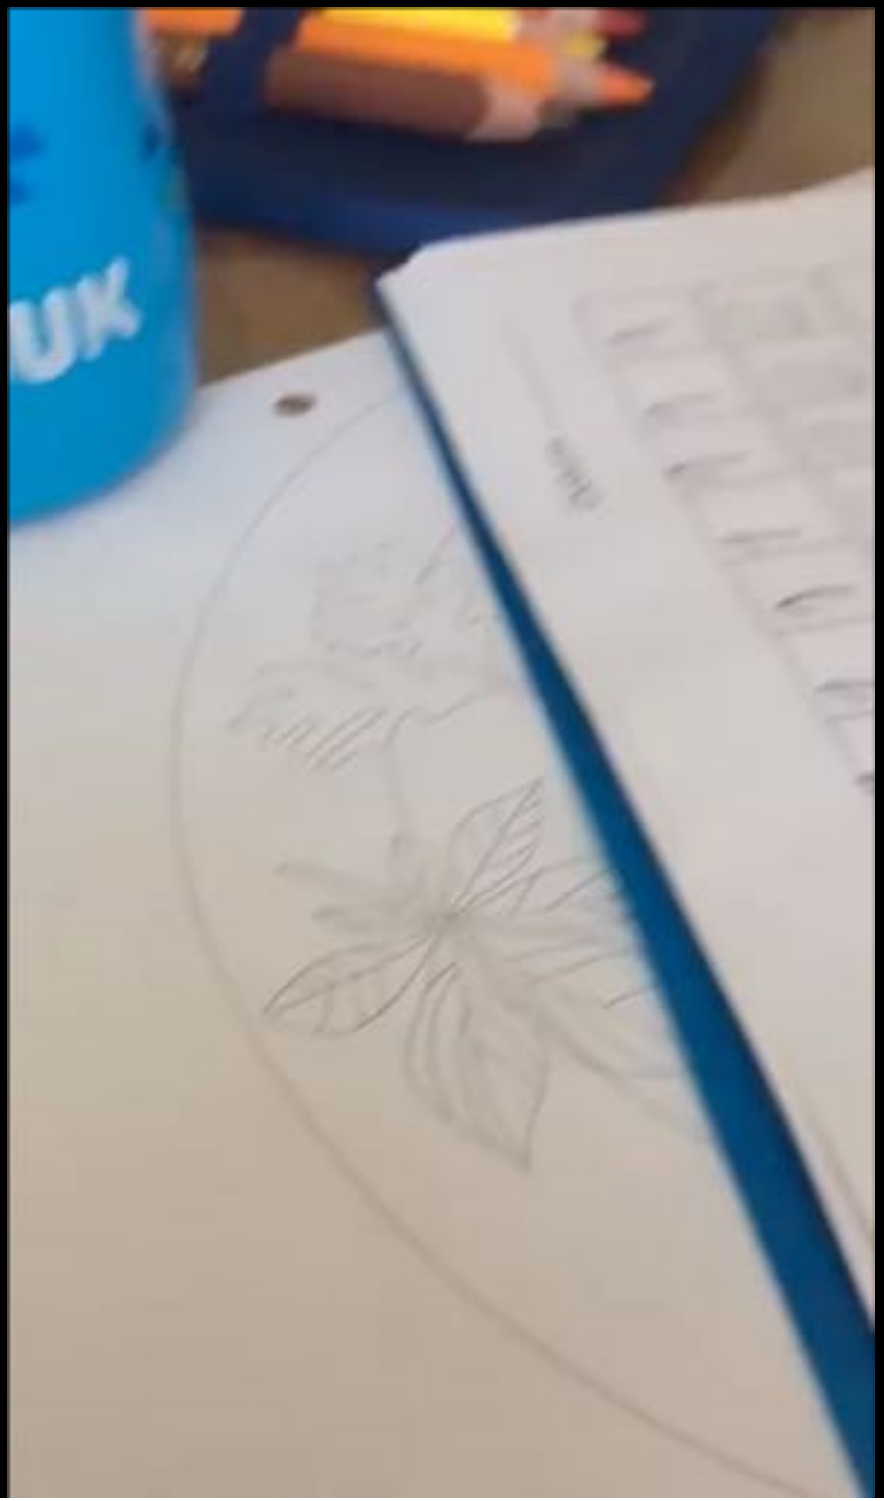

AI asks, can machines be smart, too?

„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

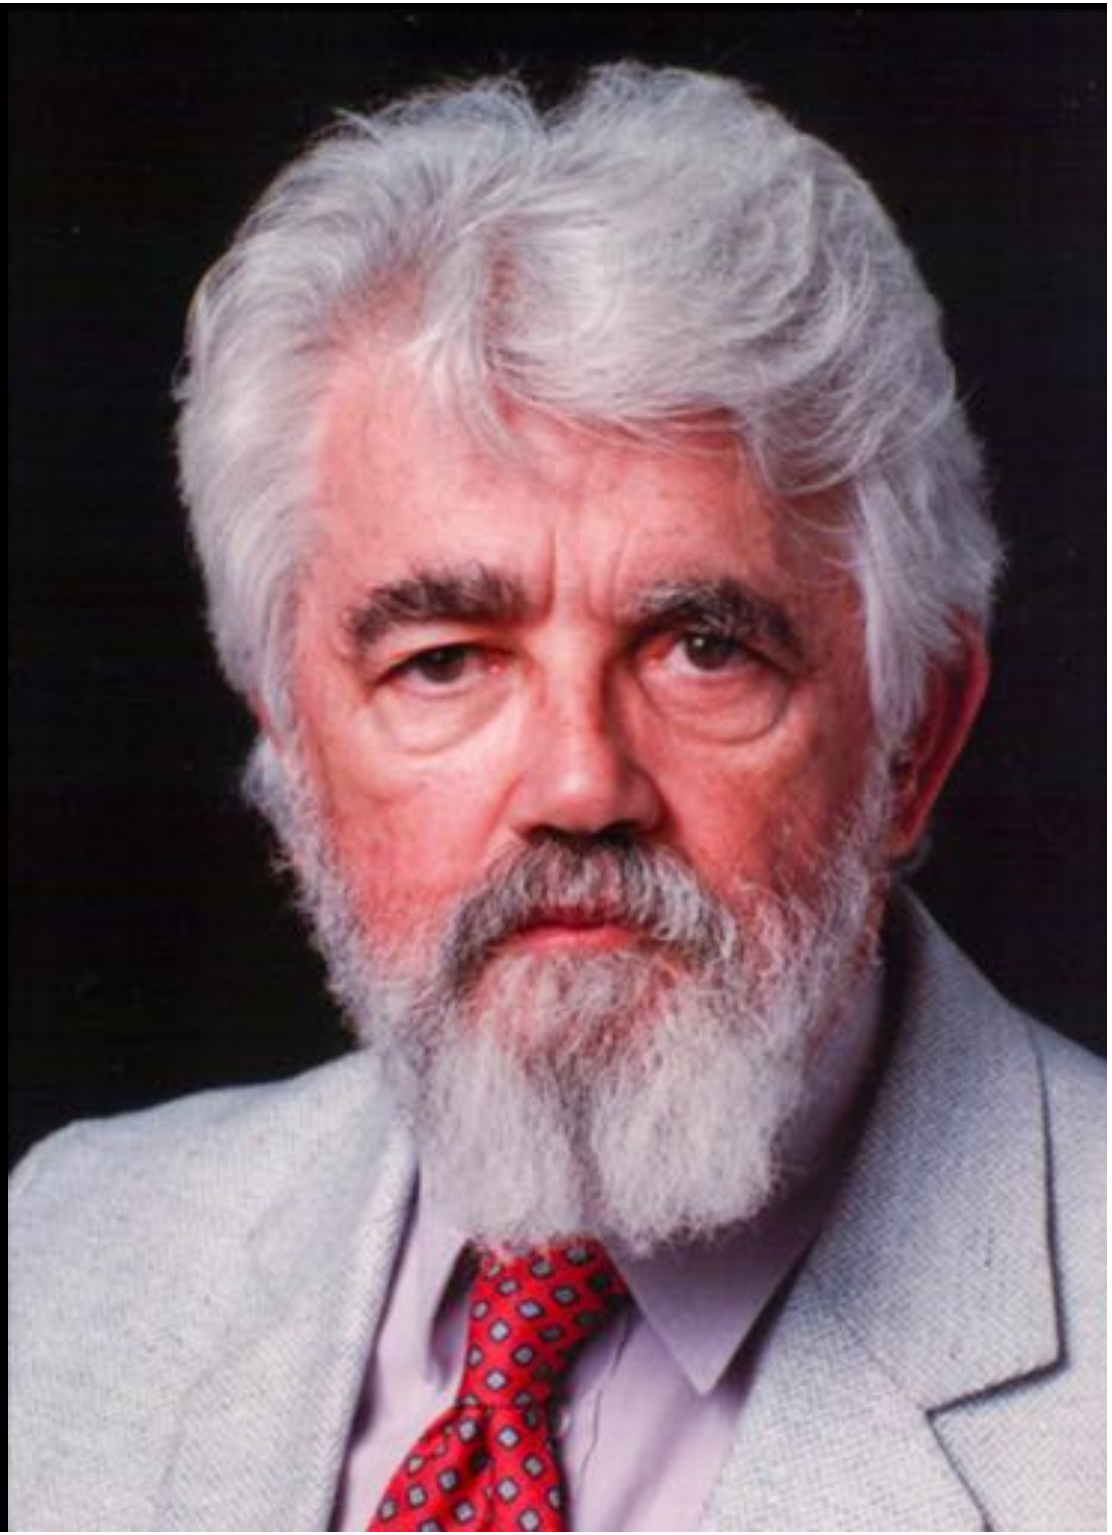

1956 – The Birth of AI

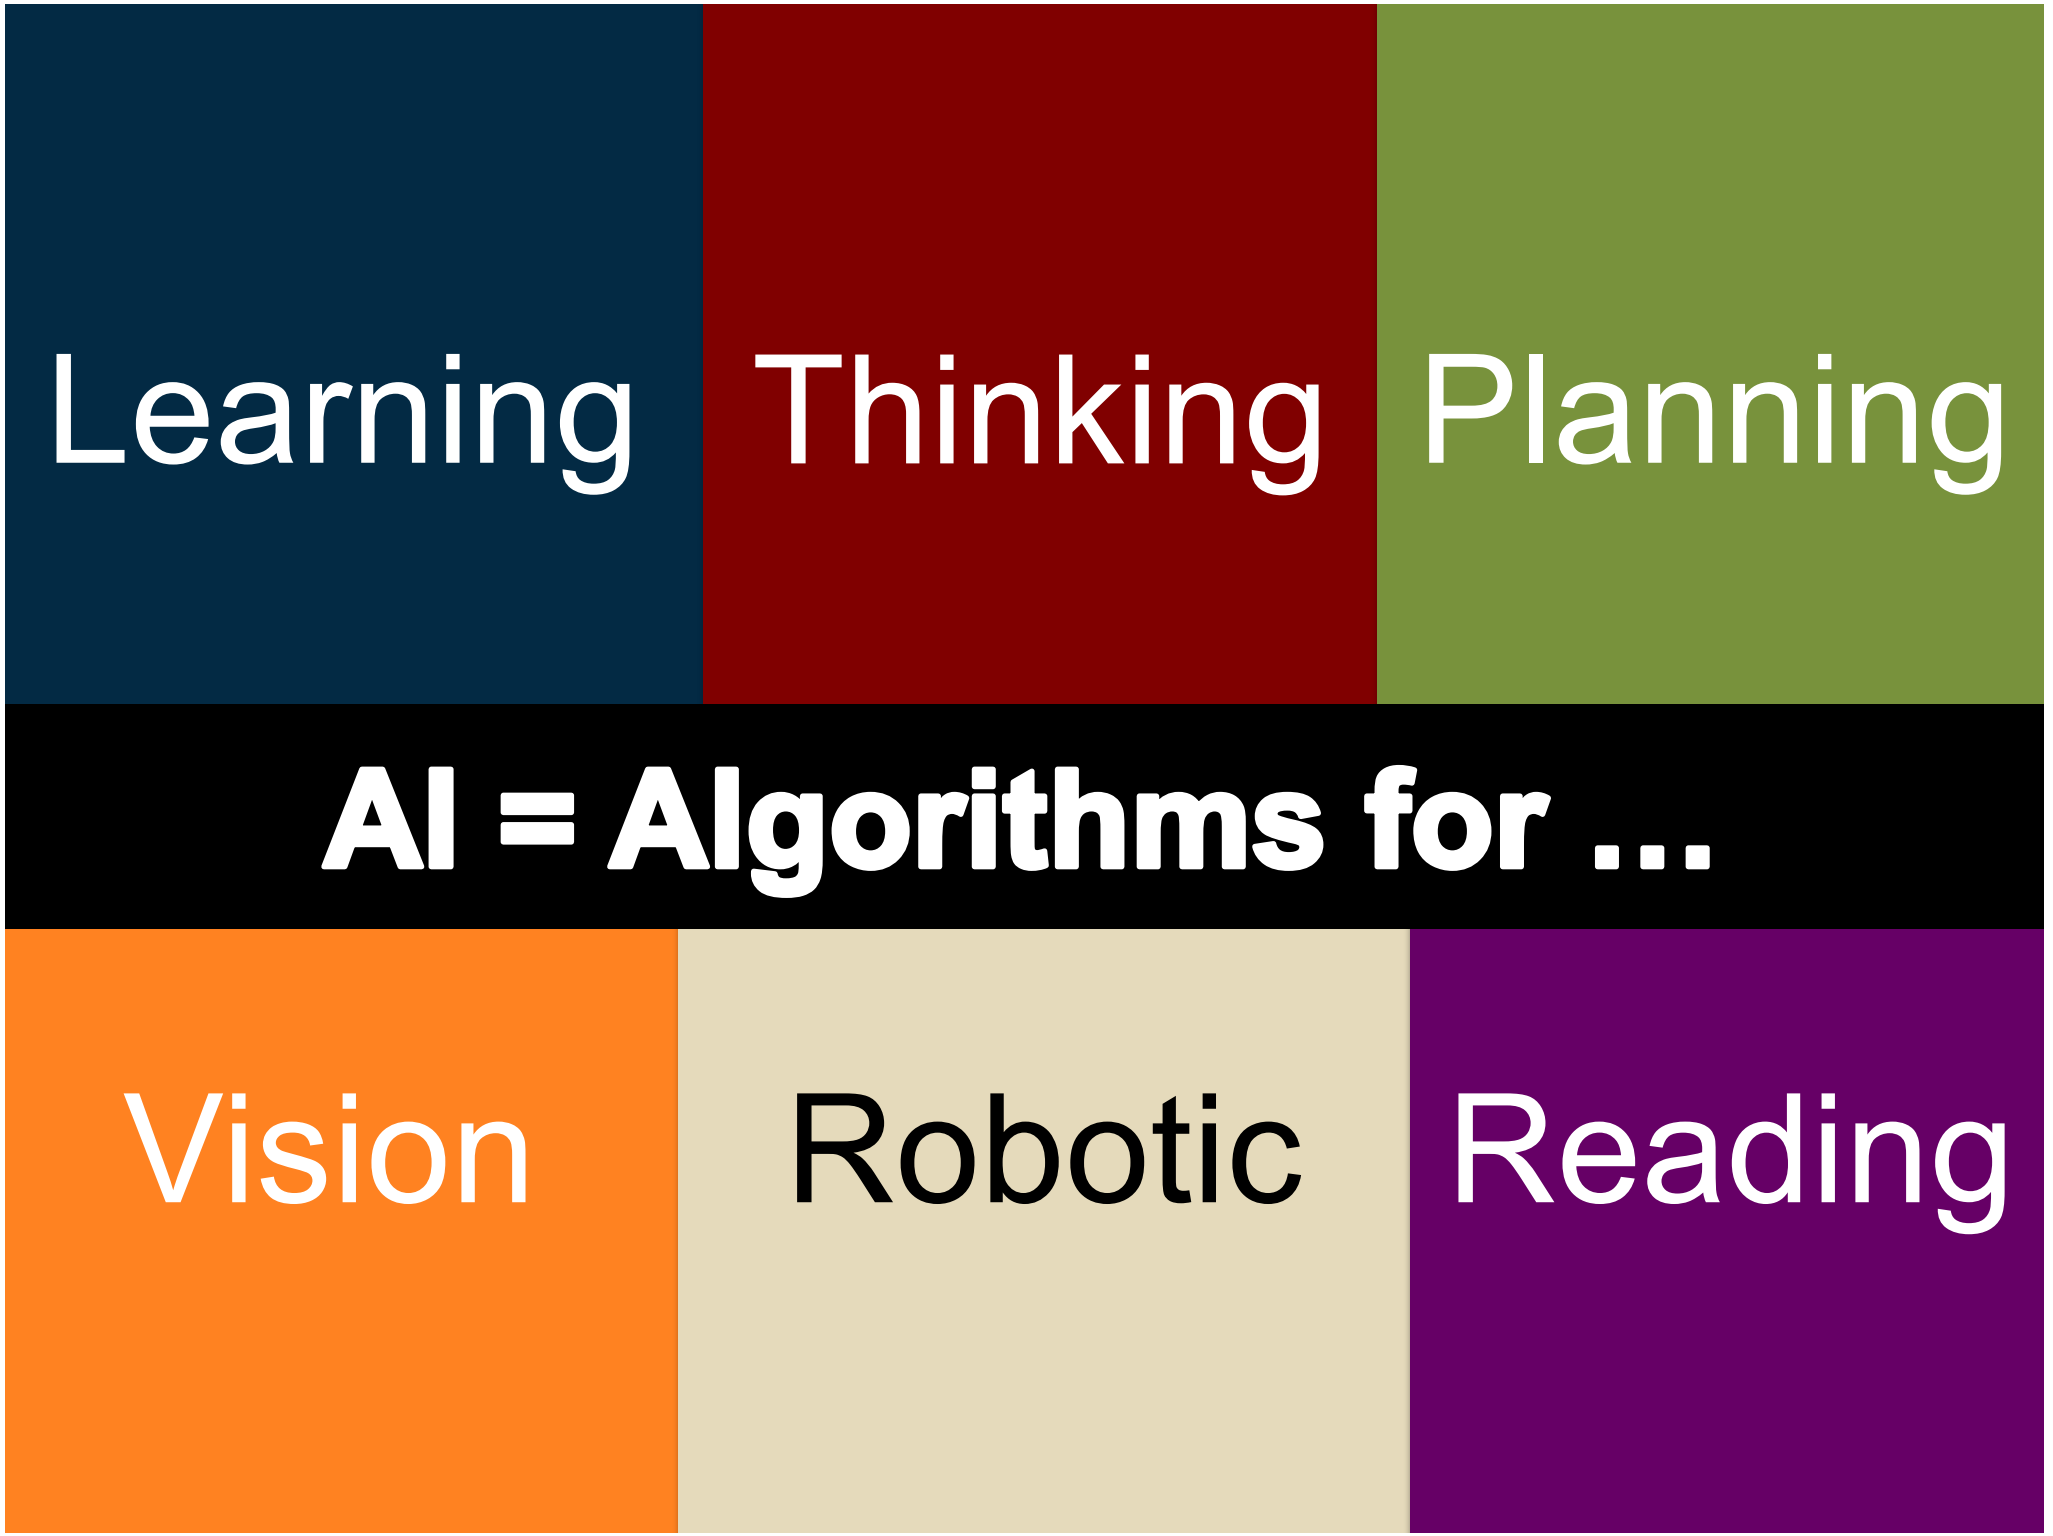
Learning
Thinking
Planning
AI = Algorithms for ...
Vision
Robotic
Reading

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)

Overall Picture

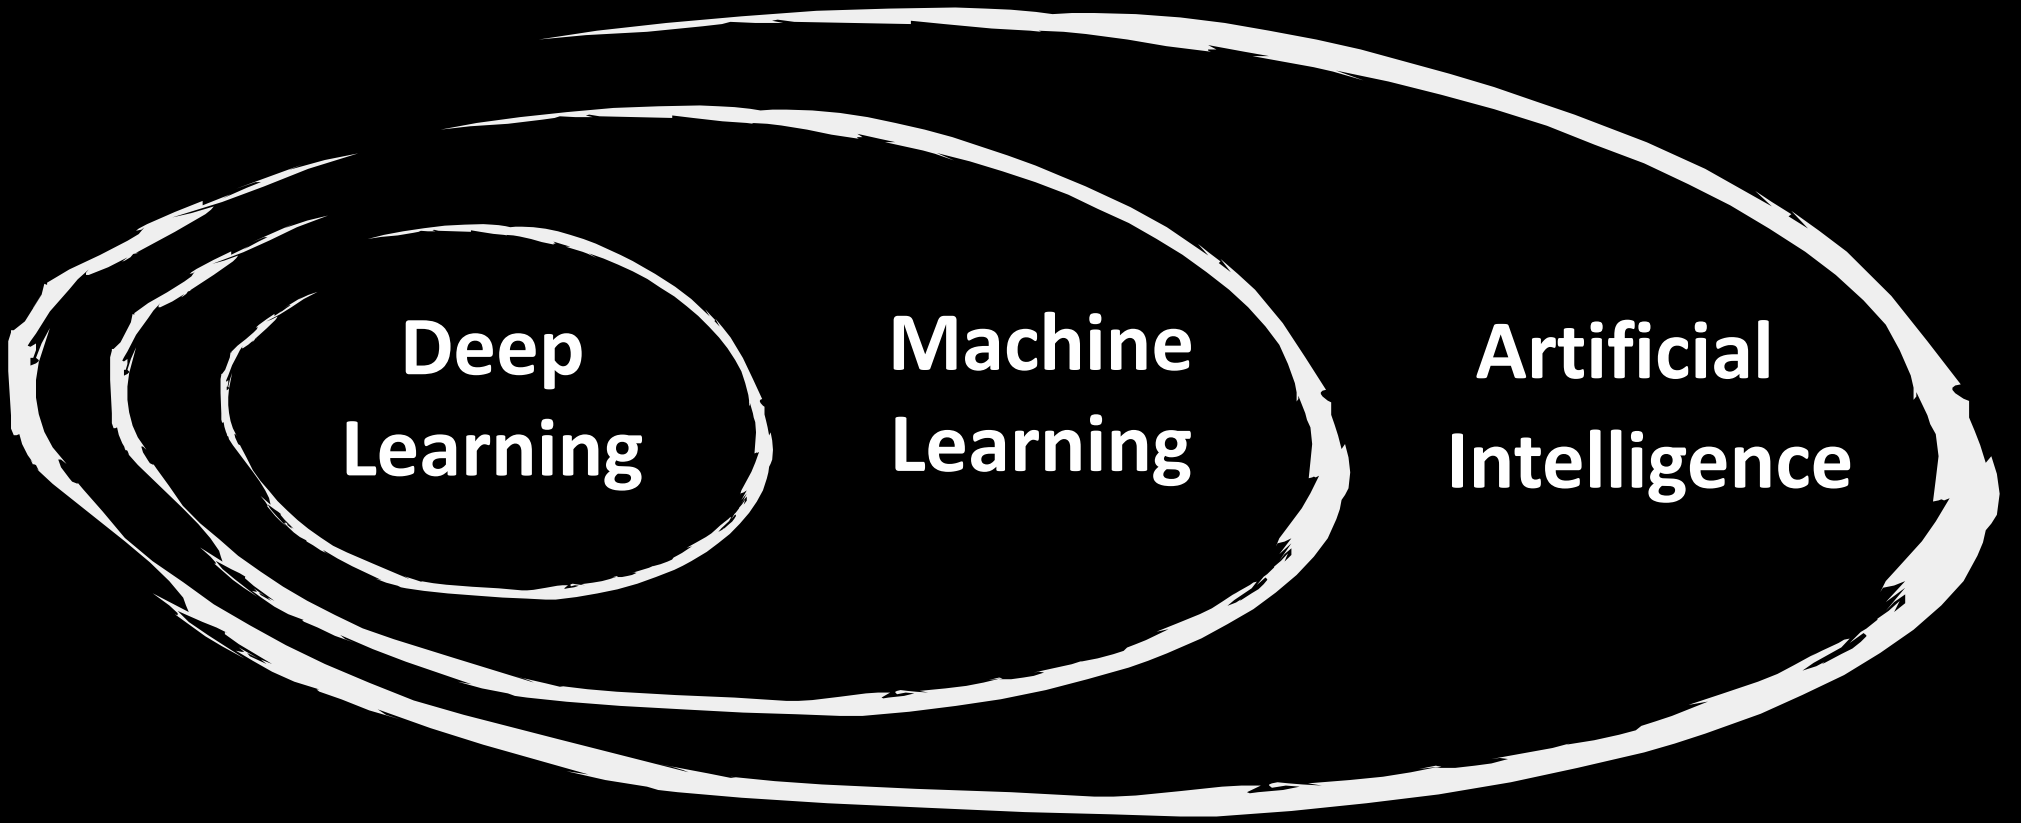


AI does the
laundry

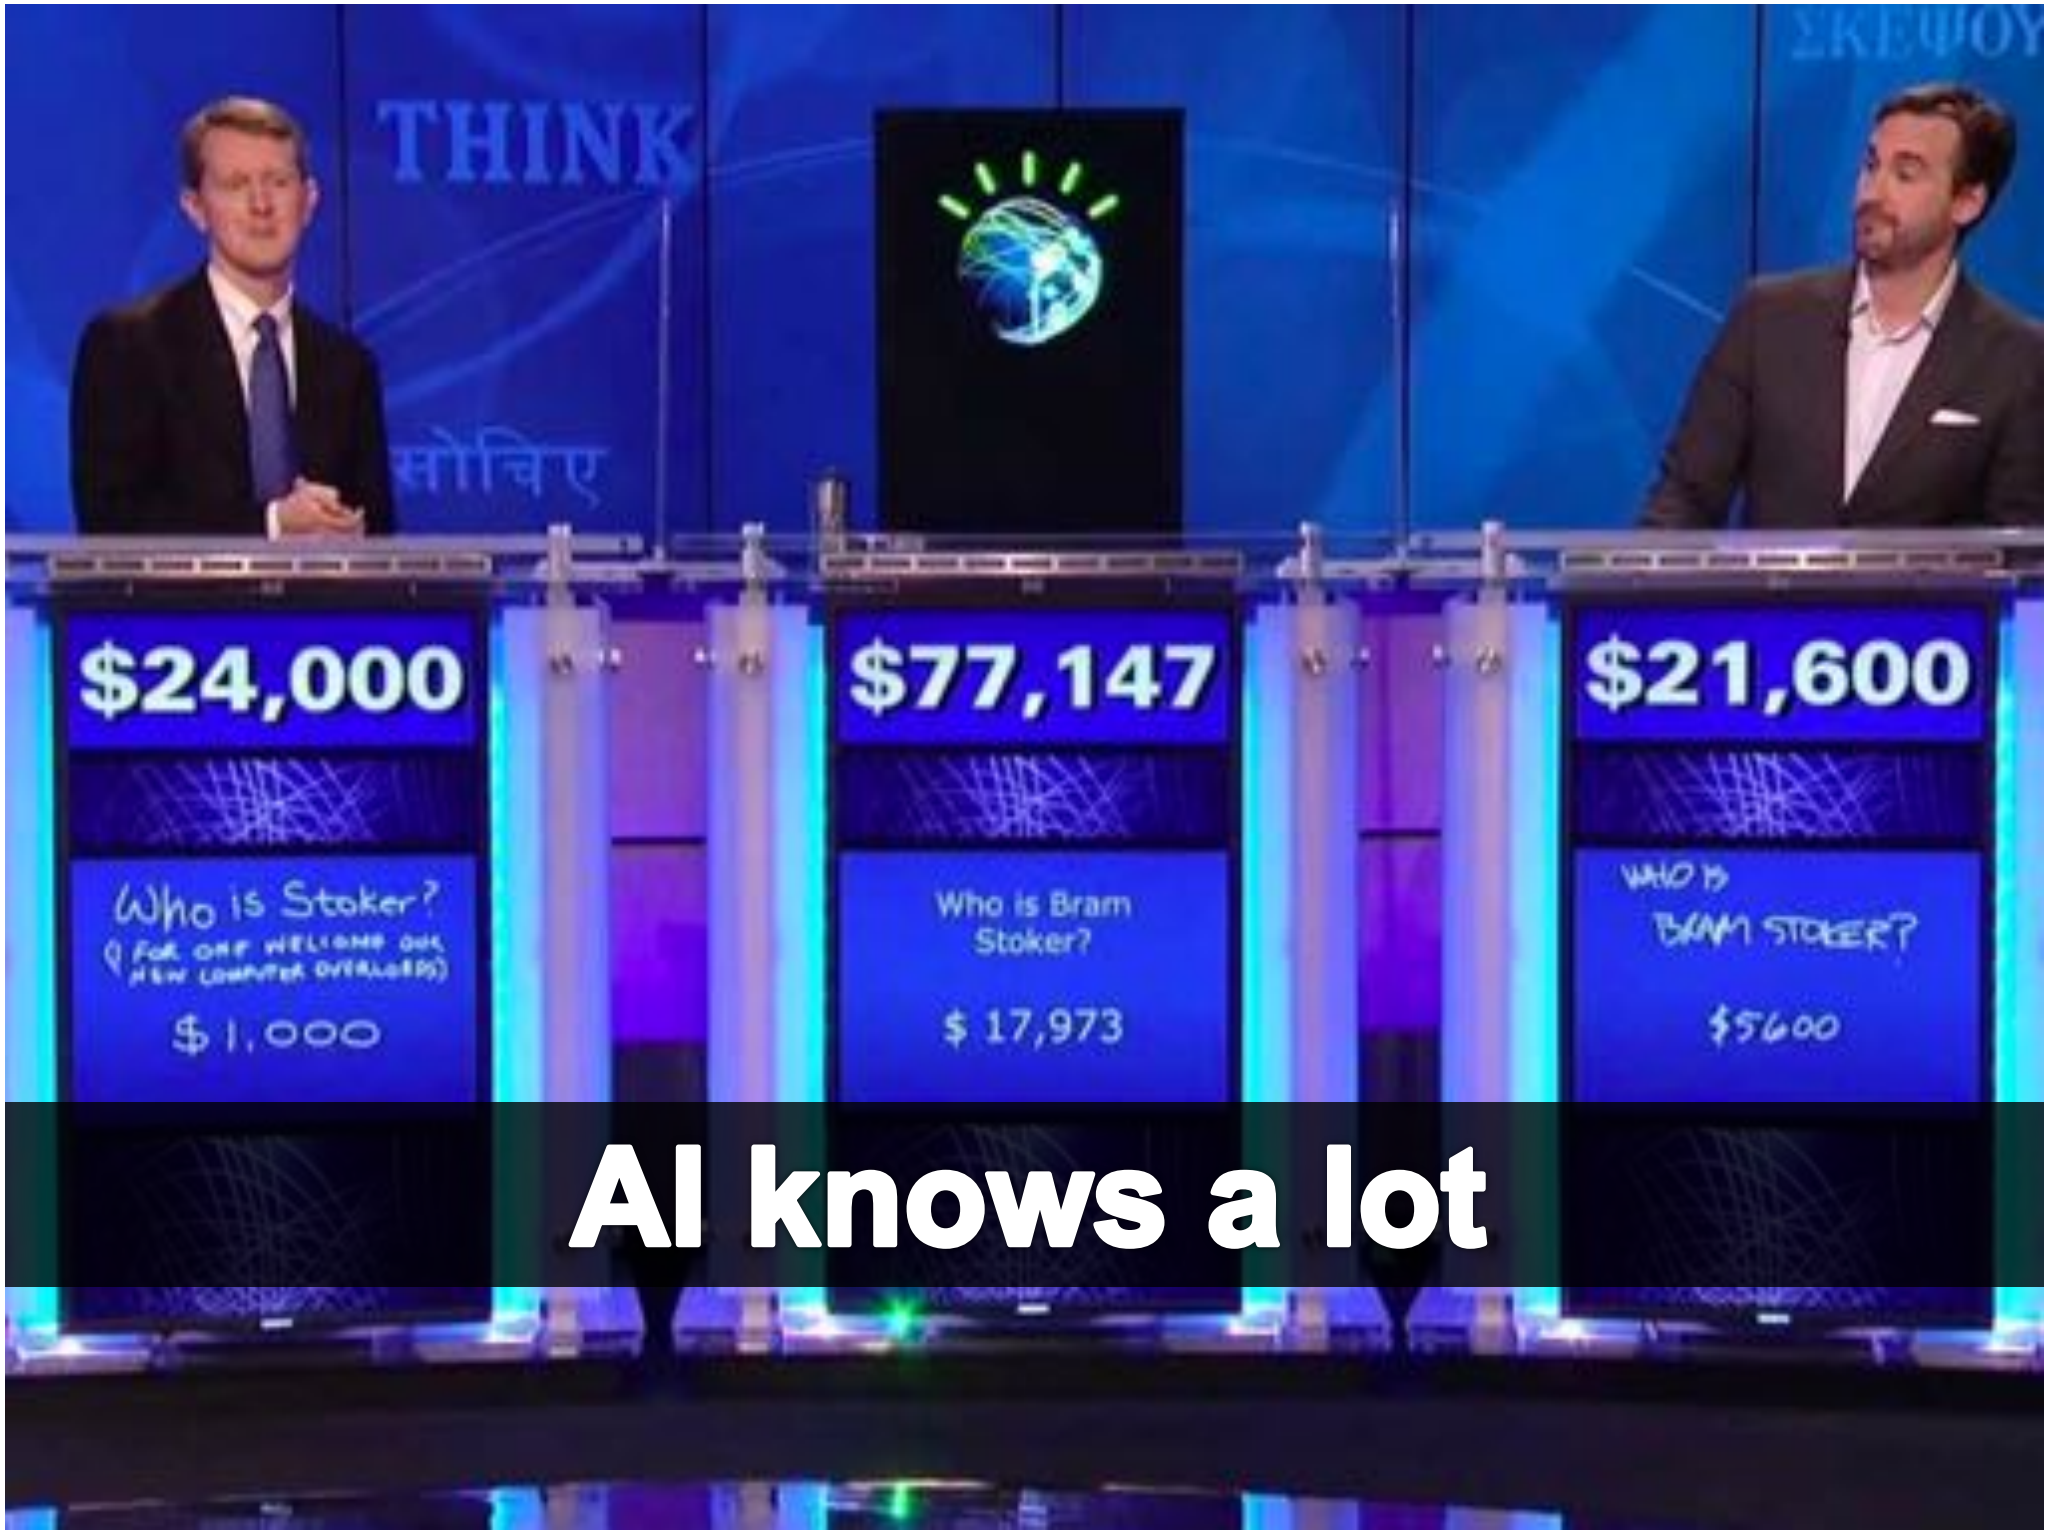
THINK
$24,000
Who is Stoker?
$ 1,000
$77,147
Who is Bram Stoker?
$ 17,973
$21,600
WHO IS BRAM STOKER?
$5600
AI knows a lot

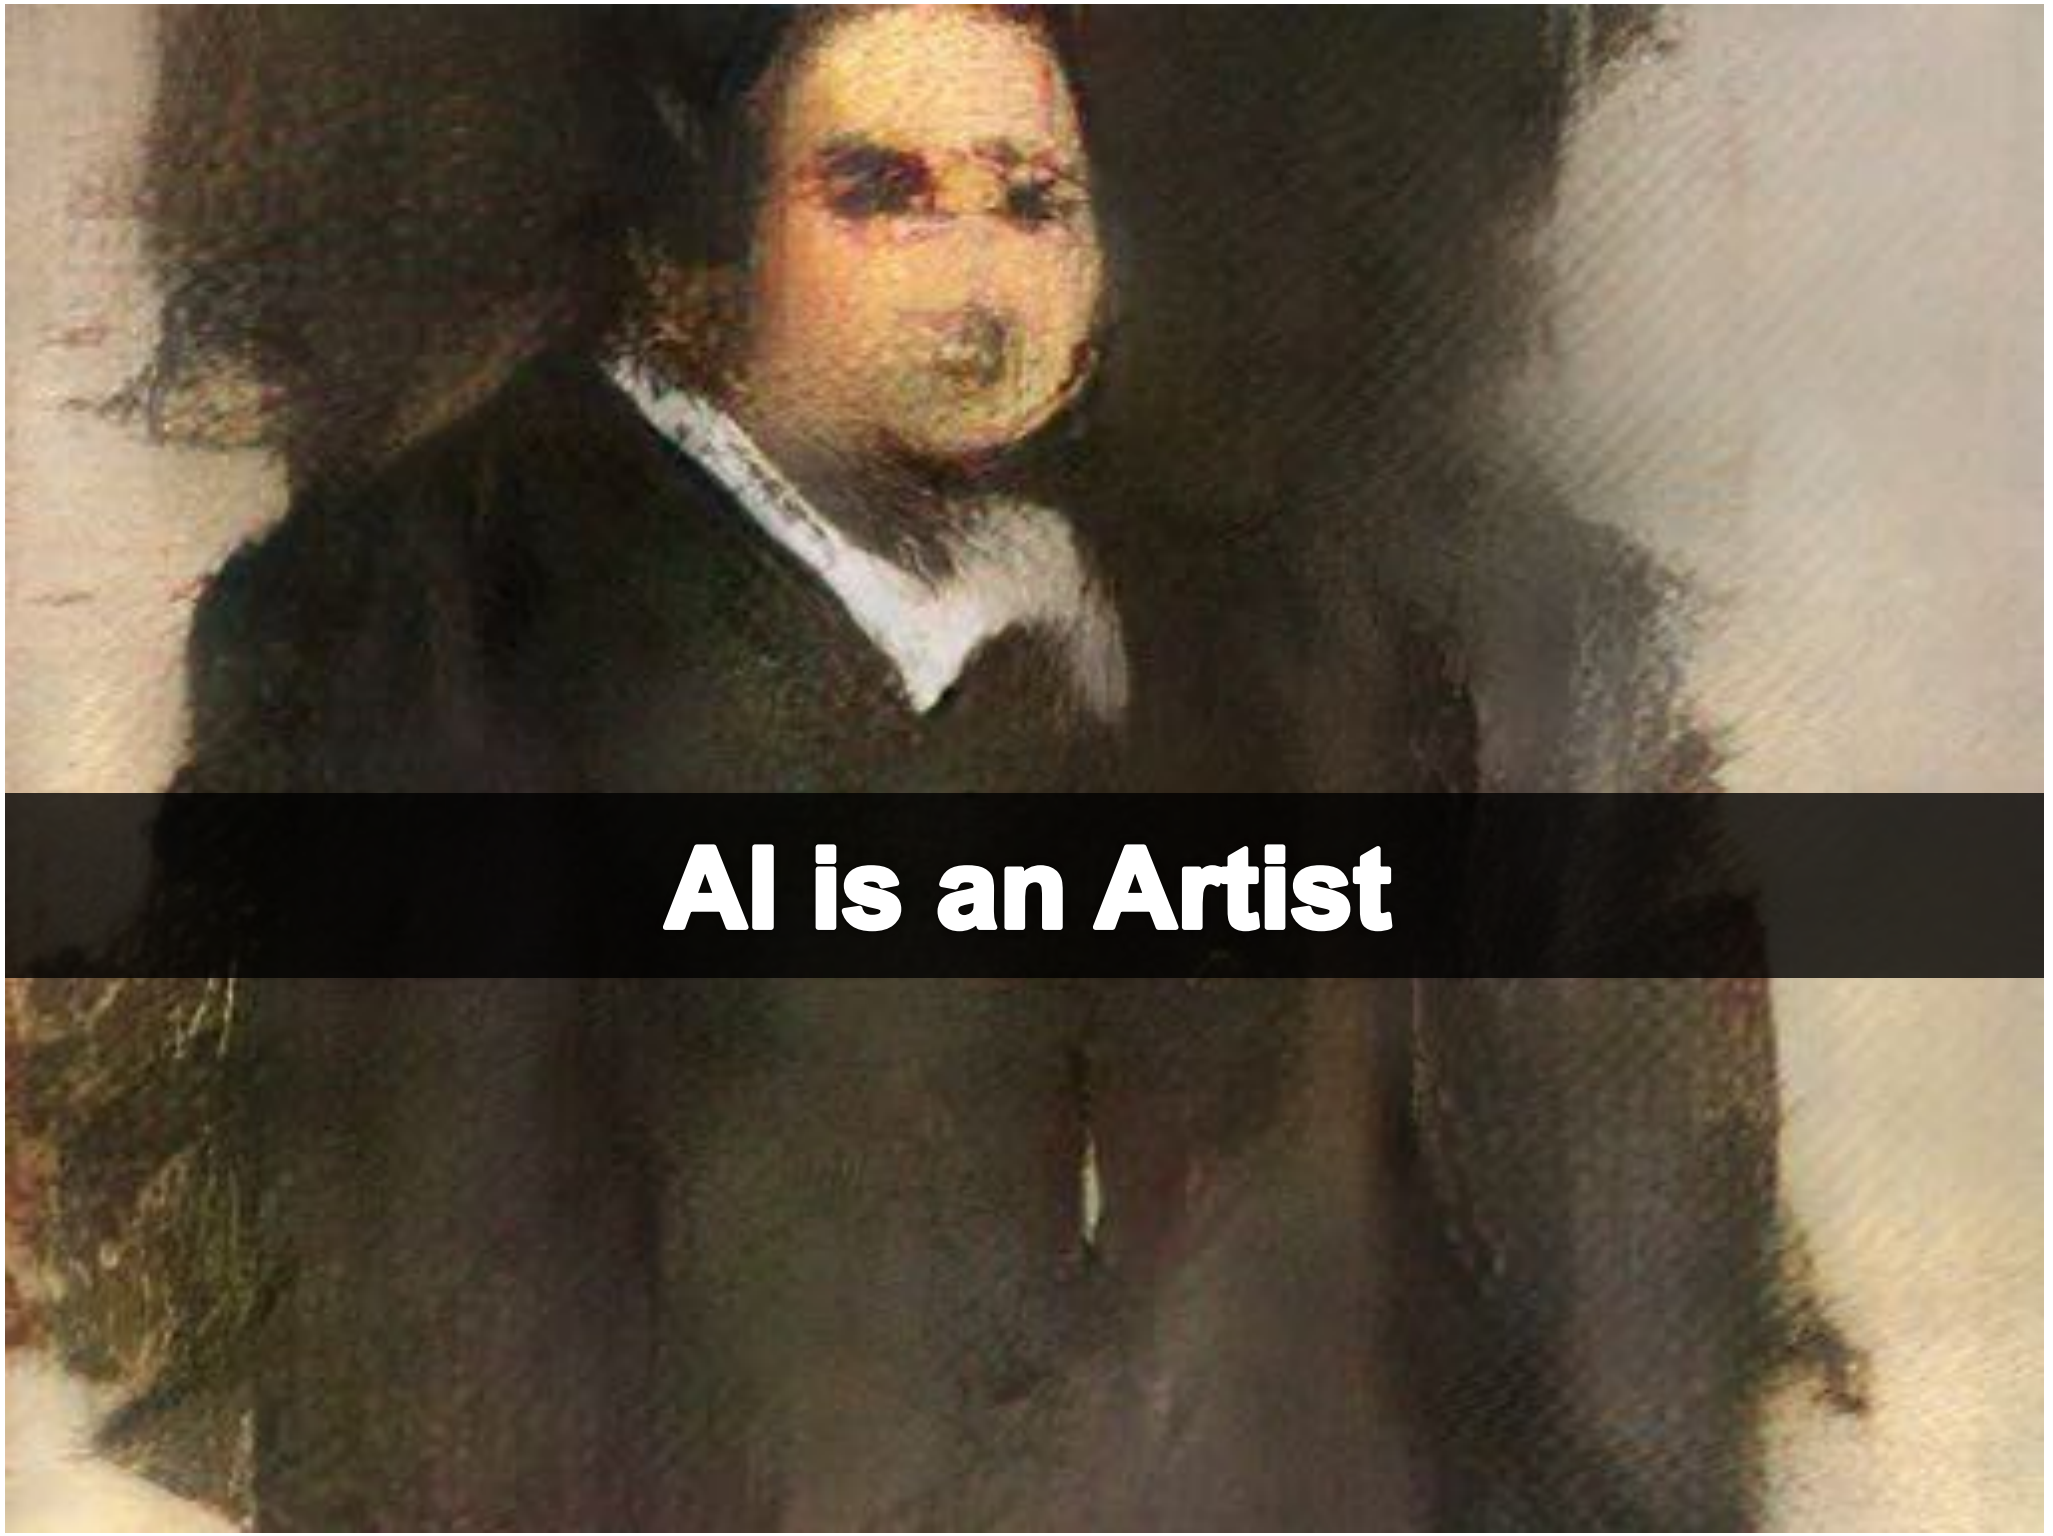
AI is an Artist


TECHNISCHE UNIVERSITÄT DARMSTADT
Schachmatt durch „CrazyAra"
19.02.2019
Künstliche Intelligenz schlägt mehrfachen Weltmeister im Einsetzschach
Der von den TU-Studierenden Johannes Czech, Moritz Willig und Alena Beyer entwickelte Bot „CrazyAra" hat den Schachprofi Justin Tan in einem Online-Match der Schach-Variante „Crazyhouse" mit 4:1 geschlagen. Gelernt hat der Bot mittels künstlicher neuronaler Netze, was ihm erlaubt, vorausschauend Entscheidungen zu treffen. Das Besondere: Die Studierenden konnten damit einen Erfolg auf einem Feld feiern, das sonst von Giganten wie Google dominiert wird.
AI plays chess and GO
lichess.org
Chat room
ALPHAGO
LEE SEDOL
01:59:31
AlphaGo
Lee Sedol
CrazyAra vs JannLee (Man vs Machine - Crazyhouse Chess on lichess.org) · 2 days ag
Category: Chess

AI assists you

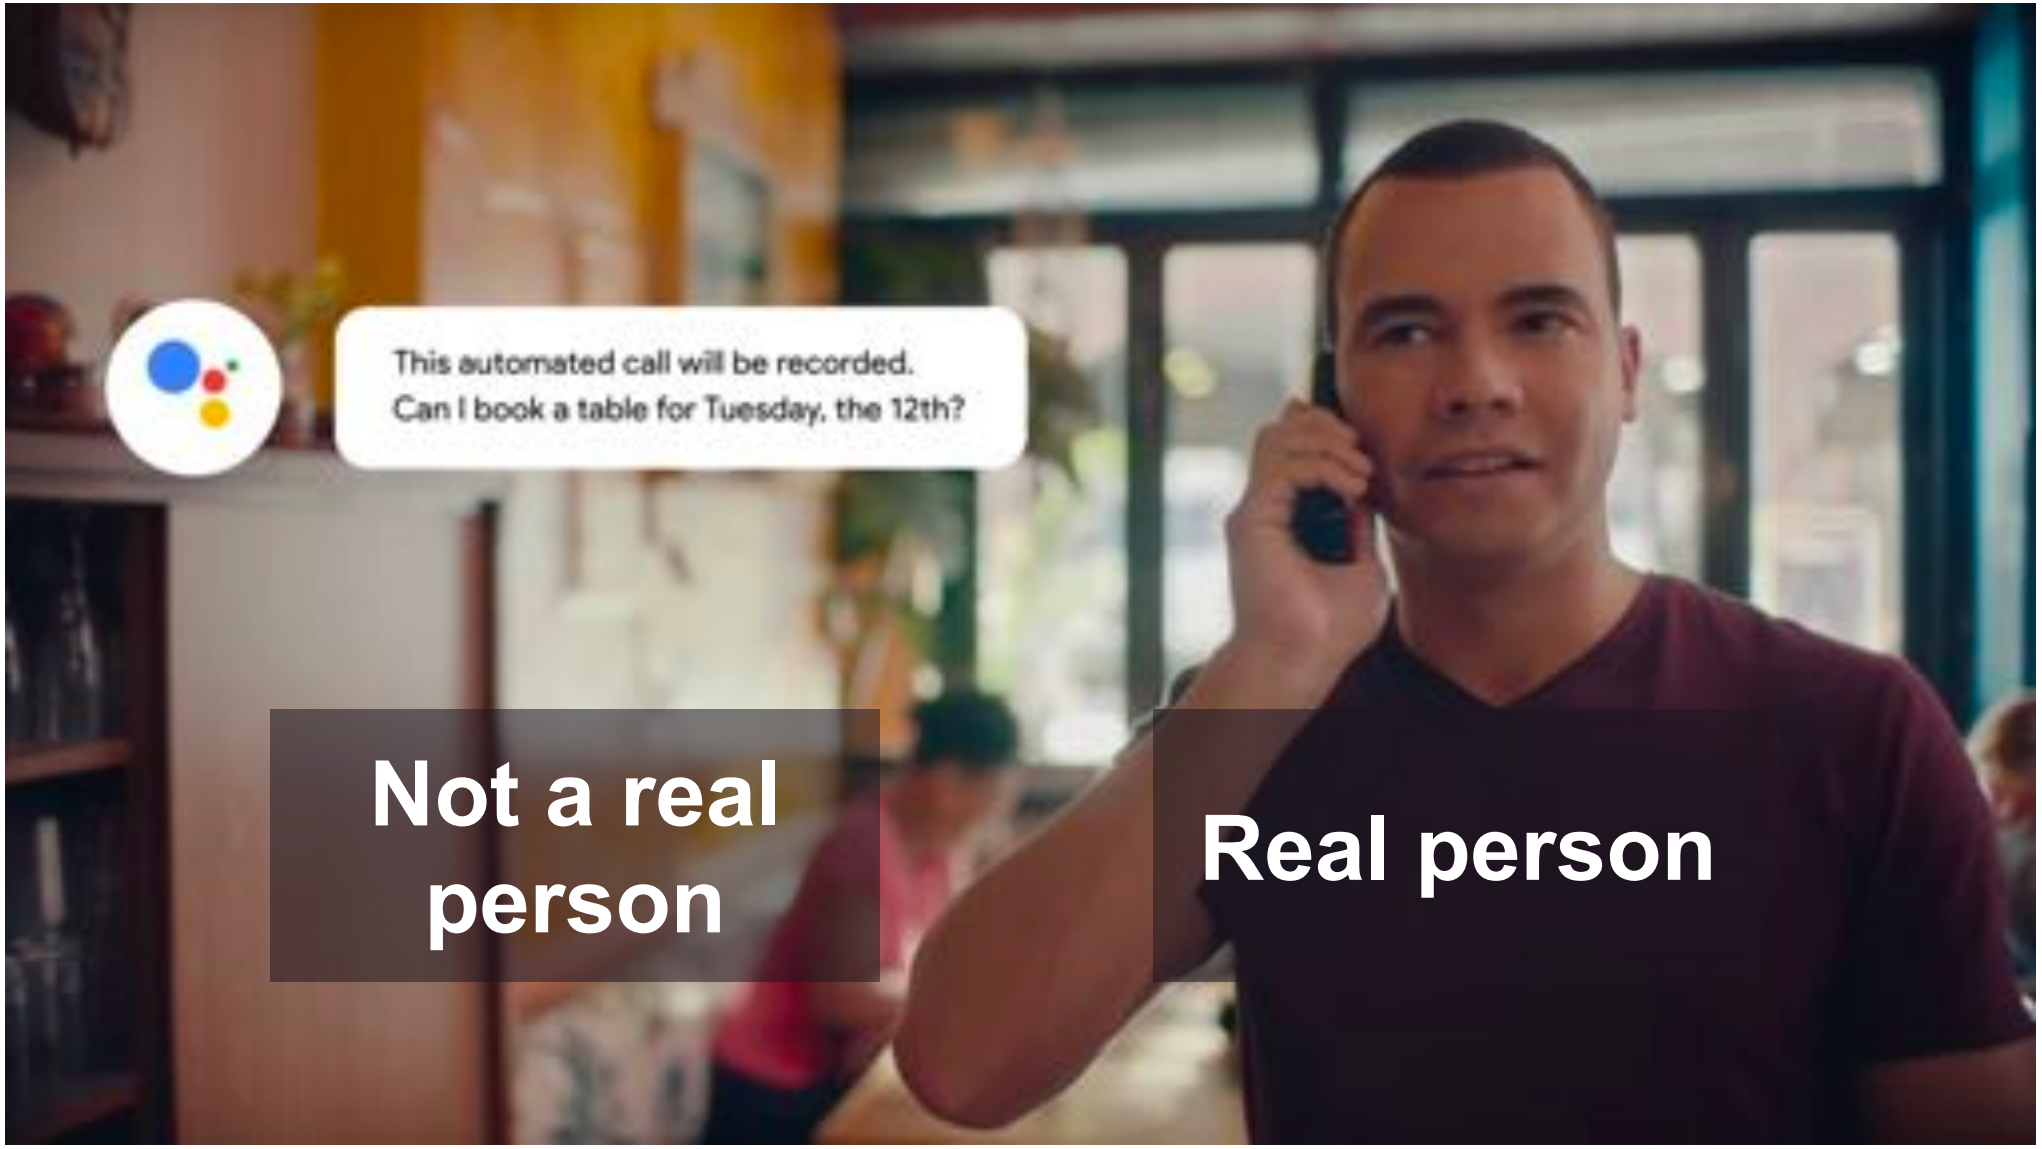

P(heart attack | )?

P(heart attack | )?

The New York Times

Opinion

A.I. Is Harder Than You Think

By Gary Marcus and Ernest Davis

Mr. Marcus is a professor of psychology and neural science. Mr. Davis is a professor of computer science.

May 18, 2018

P(heart attack | )?

The New York Times

Opinion

A.I. Is Harder Than You Think

By Gary Marcus and Ernest Davis

Mr. Marcus is a professor of psychology and neural science. Mr. Davis is a professor of computer science.

May 18, 2018

Standard machine learning

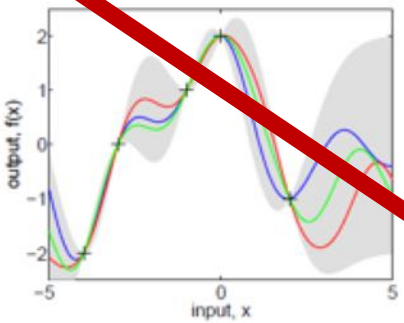

Gaussian Processes

Graphical models

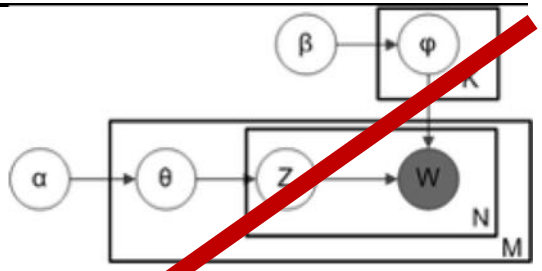

Features

Objects

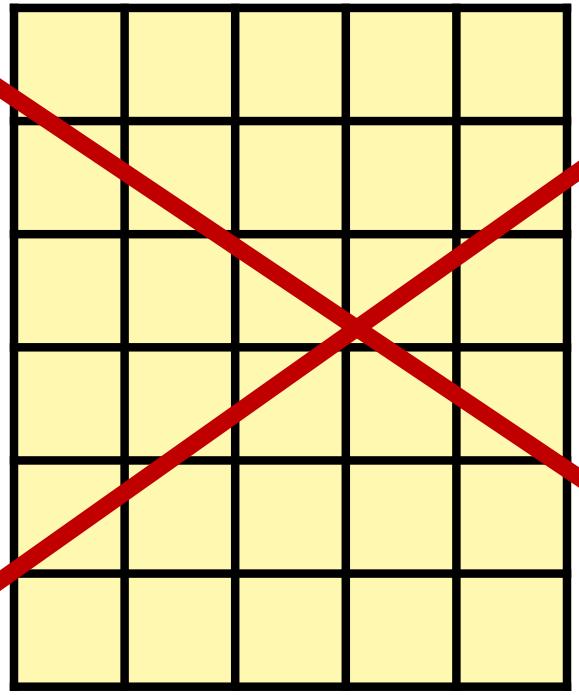

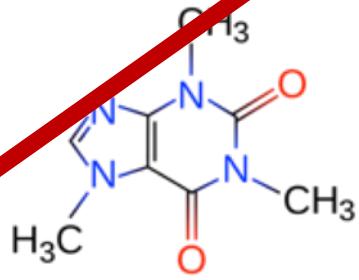

Graph Mining

Big Model

teaches

Small Model

Distillation/LUPI

Boosting

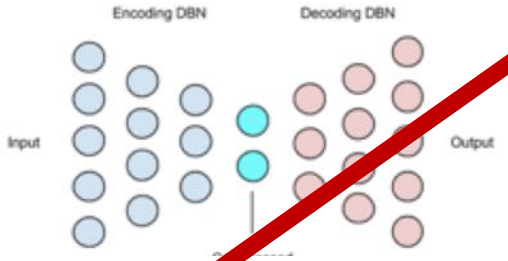

Autoencoder,
Deep Learning

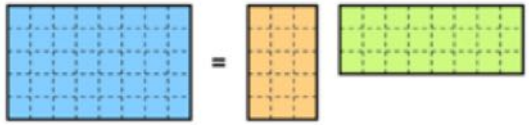

Big Data Matrix Factorization

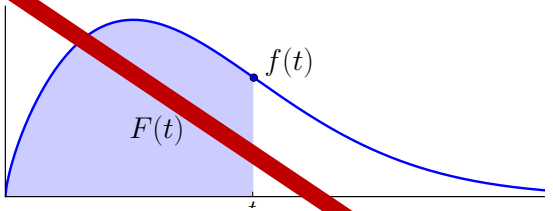

Diffusion Models

and many more …

EHRs are dirty and interconnected

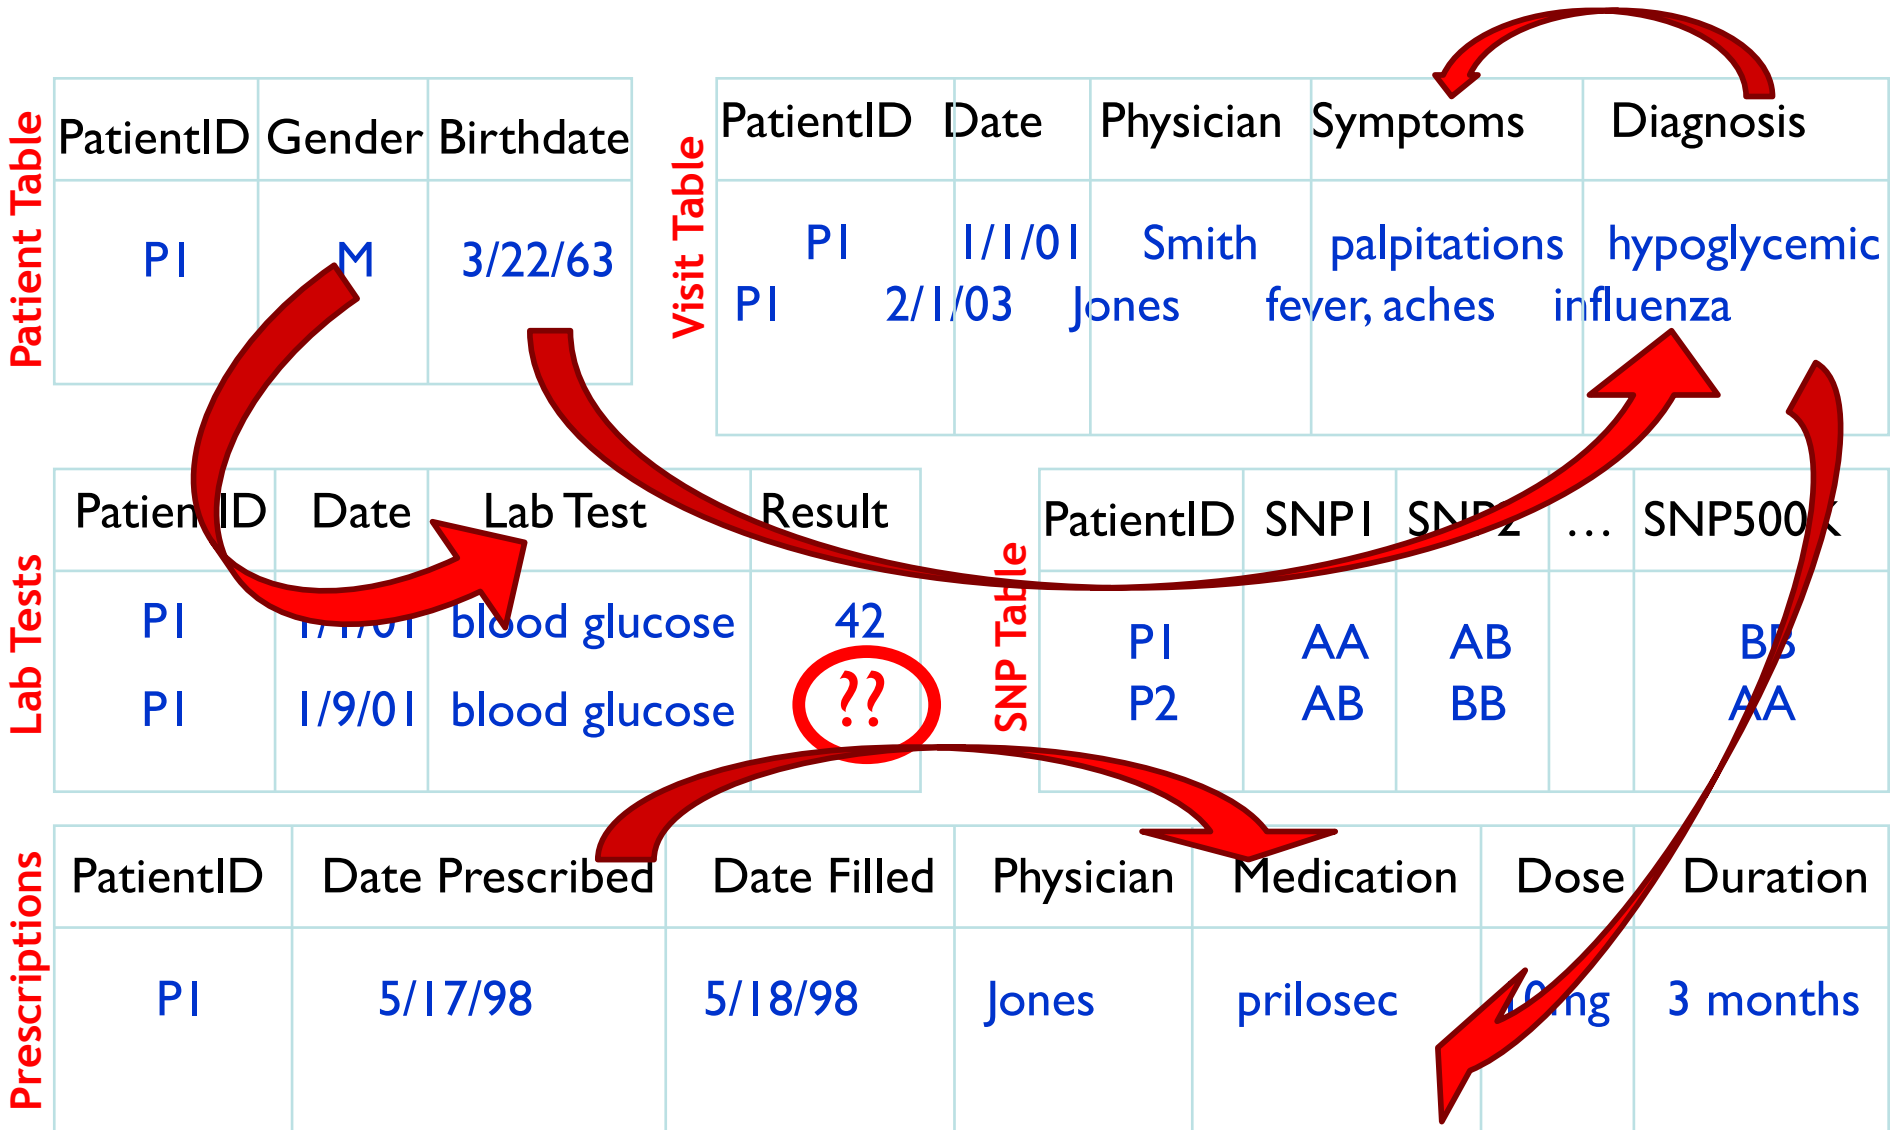

Natarajan, Khot, Kersting, Shavlik. Boosted Statistical Relational Learners. Springer Brief 2015

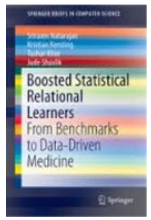

Understanding Electronic Health Records

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

TECHNISCHE UNIVERSITÄT DARMSTADT

THE UNIVERSITY OF TEXAS AT DALLAS

Logical Variables
(Abstraction)

Rule/Database view

Probability

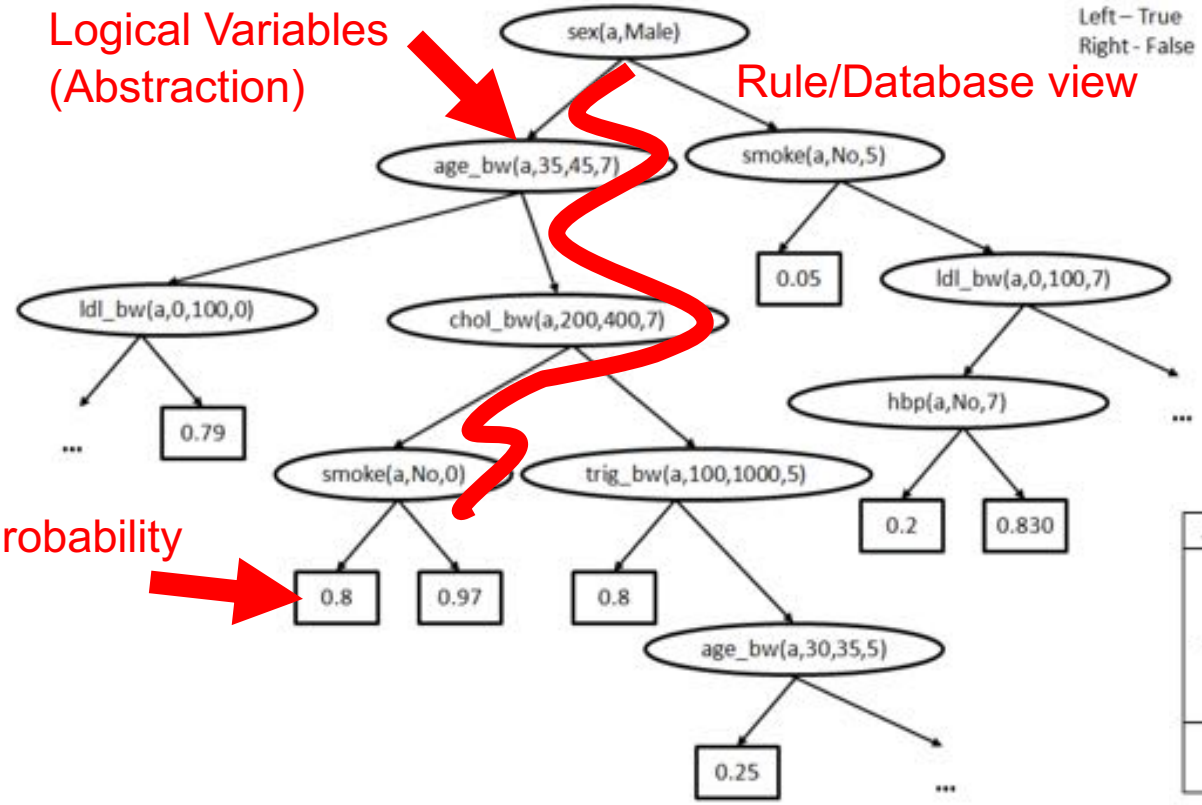

[Circulation; 92(8), 2157-62, 1995; JACC; 43, 842-7, 2004]

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]


https://www.youtube.com/watch?v=sdUHX72qxeY
1863

Current Biology

Search
All Content
Advanced Search
Current Biology All Journals

Explore Online Now Current Issue Archive Journal Information For Authors

< Previous Article
Volume 27, Issue 18, p2827–2832.e3, 25 September 2017
Next Article >

REPORT

Humans, but Not Deep Neural Networks, Often Miss Giant Targets in Scenes

Miguel P. Eckstein, Kathryn Koehler, Lauren E. Welbourne, Emre Akbas

Optical Illusions

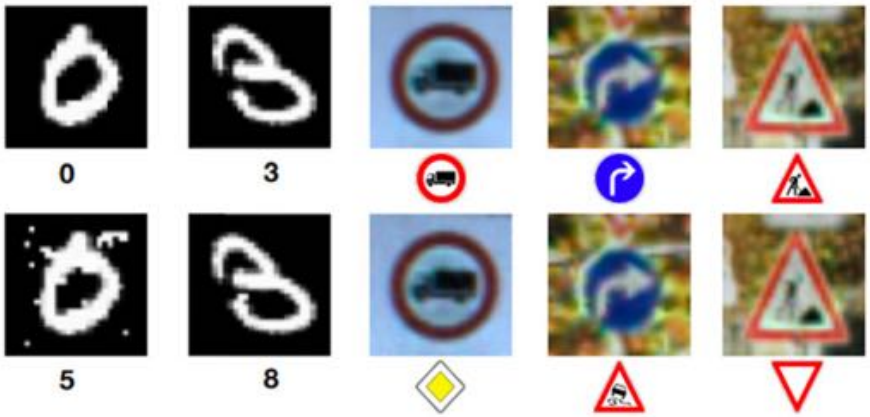

Stereotypes

Target words	Attrib. words	Original Finding				Our Finding			
		Ref	N	d	p	N_T	N_A	d	p
Flowers vs insects	Pleasant vs unpleasant	(5)	32	1.35	10^{-8}	25×2	25×2	1.54	10^{-7}
Instruments vs weapons	Pleasant vs unpleasant	(5)	32	1.66	10^{-10}	25×2	25×2	1.63	10^{-8}
Eur.-American vs Afr.-American names	Pleasant vs unpleasant	(5)	26	1.17	10^{-5}	32×2	25×2	0.58	10^{-2}
Eur.-American vs Afr.-American names	Pleasant vs unpleasant	(7)	Not applicable			18×2	25×2	1.24	10^{-3}
Eur.-American vs Afr.-American names	Pleasant vs unpleasant from (5)	(7)	Not applicable			18×2	8 × 2	0.72	10^{-2}
Male vs female names	Career vs family	(9)	39k	0.72	10^{-2}	8 × 2	8 × 2	1.89	10^{-4}
Math vs arts	Male vs female terms	(9)	28k	0.82	$< 10^{-2}$	8 × 2	8 × 2	0.97	.027
Science vs arts	Male vs female terms	(10)	91	1.47	10^{-24}	8 × 2	8 × 2	1.24	10^{-2}
Mental vs physical disease	Temporary vs permanent	(23)	135	1.01	10^{-3}	6 × 2	7 × 2	1.30	.012
Young vs old people's names	Pleasant vs unpleasant	(9)	43k	1.42	$< 10^{-2}$	8 × 2	8 × 2	−.08	0.57

The third wave of AI

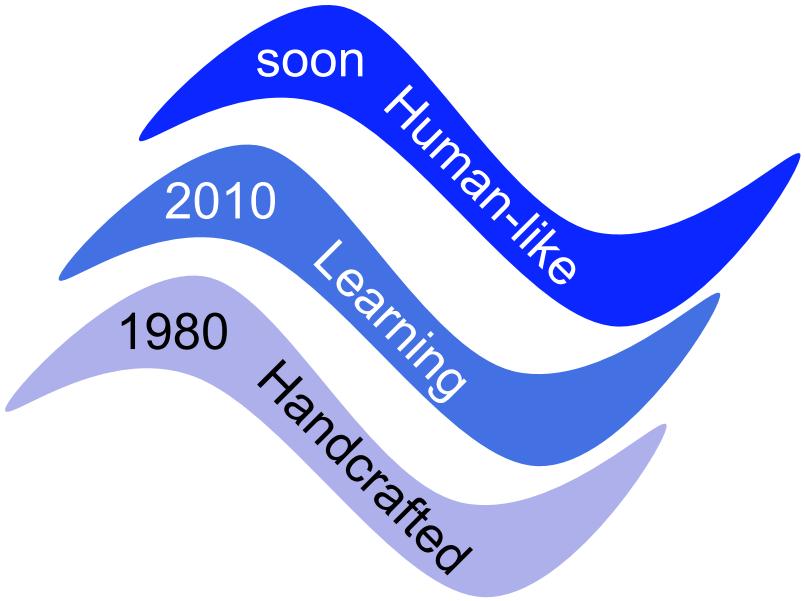

AI systems that can acquire human-like communication and reasoning capabilities, with the ability to recognise new situations and adapt to them.

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]

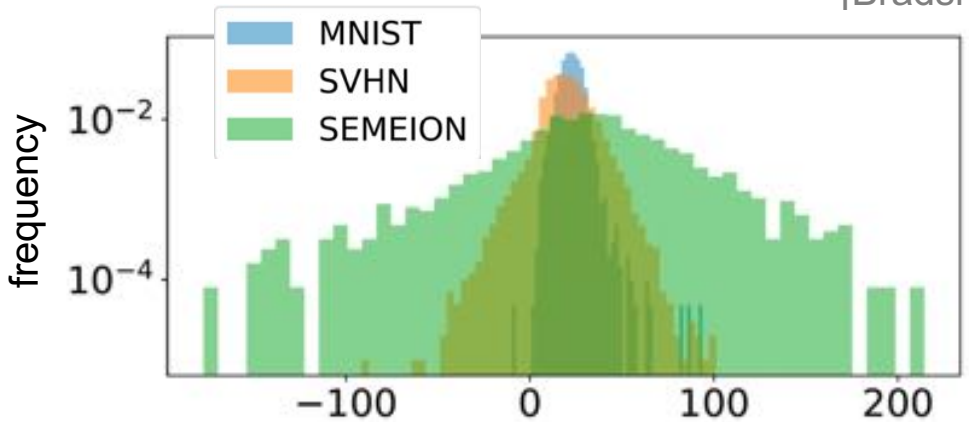

Many DNNs cannot distinguish the datasets

[Peharz, Vergari, Molina, Stelzner, Trapp, Kersting, Ghahramani UDL@UAI 2018]

Random sum-product networks

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

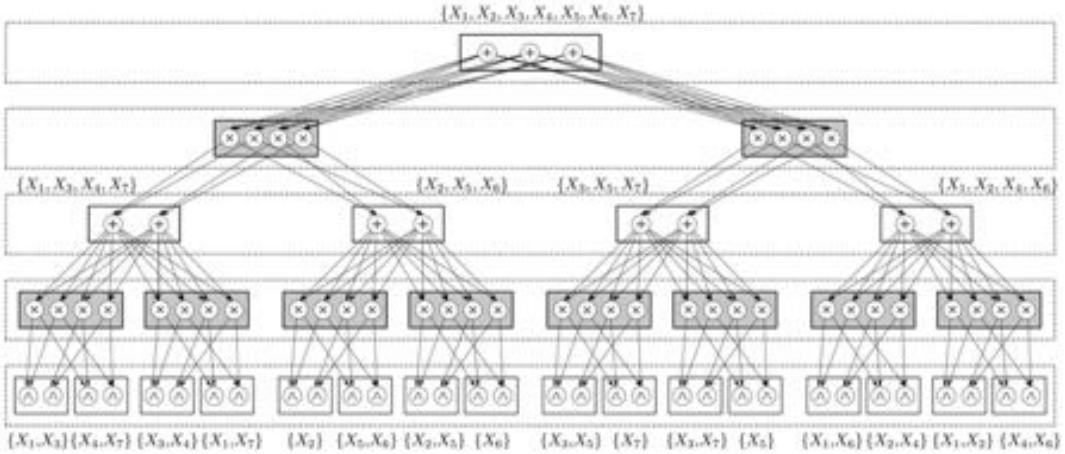

Build a random SPN structure. This can be done in an informed way or completely at random

		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 (9.38M)	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)

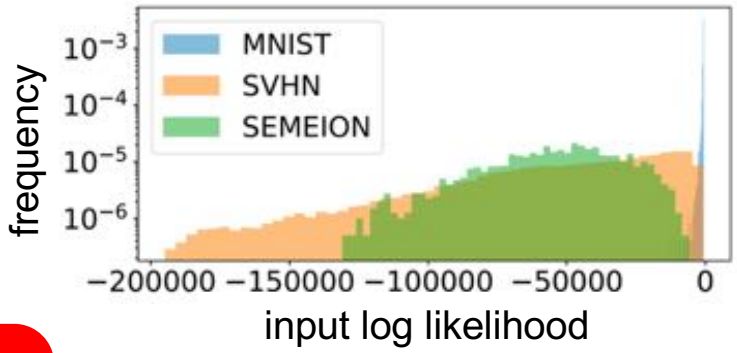

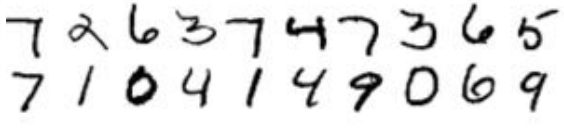

outliers

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

[Sommer, Oppermann, Molina, Binnig, Kersting, Koch ICDD 2018]

TABLE II
PERFORMANCE COMPARISON. BEST END-TO-END THROUGHPUTS (T), EXCLUDING THE CYCLE COUNTER MEASUREMENTS, ARE DENOTED BOLD.

Dataset	Rows	CPU (µs)	T-CPU (rows/µs)	CPUF (µs)	T-CPUF (rows/µs)	GPU (µs)	T-GPU (rows/µs)	FPGA Cycle Counter	FPGAC (µs)	T-FPGAC (rows/µs)	FPGA (µs)	T-FPGA (rows/µs)
Accidents	17009	2798.27			7.87	63090.94	0.27	17249			696.00	**24.44**
Audio	20000	4271.78			5.4			20317	1		761.00	**26.28**
Netflix	20000	4892.22			4.8			20322	1		654.00	**30.58**
MSNBC200	388434	15476.05			30.5			388900	19		008.00	**77.56**
MSNBC300	388434	10060.78			41.2			388810	19		933.00	**78.74**
NLTCS	21574	791.80			31.3			21904	1		566.00	**38.12**
Plants	23215	3621.71	6.41	3521.04	6.59	67004.41	0.35	23592	117.96	196.80	778.00	**29.84**
NIPS5	10000	25.11	**398.31**	26.37	379.23	8210.32	1.22	10236	51.18	195.39	337.30	29.65
NIPS10	10000	83.60	**119.61**	84.39	118.49	11550.82	0.87	10279	51.40	194.57	464.30	21.54
NIPS20	10000	191.30	52.27	182.73	**54.72**	18689.04	0.54	10285	51.43	194.46	543.60	18.40
NIPS30	10000	387.61	25.80	349.84	**28.58**	25355.93	0.39	10308	51.80	193.06	592.30	16.88
NIPS40	10000	551.64	18.13	471.26	**21.22**	30820.49	0.32	10306	51.53	194.06	632.20	15.82
NIPS50	10000	812.44	12.31	792.13	12.62	36355.60	0.28	10559	52.80	189.41	720.60	**13.88**
NIPS60	10000	1046.38	9.56	662.53	**15.09**	40778.36	0.25	12271	61.36	162.99	799.20	12.51
NIPS70	10000	1148.17	8.71	1134.80	8.81	46759.26	0.21	14022	70.11	142.63	858.60	**11.65**
NIPS80	10000	1556.99	6.42	1277.81	7.83	63217.99	0.16	14275	78.51	127.37	961.80	**10.40**

How do we do data science offshore?

MADESI

Federal Ministry of Education and Research

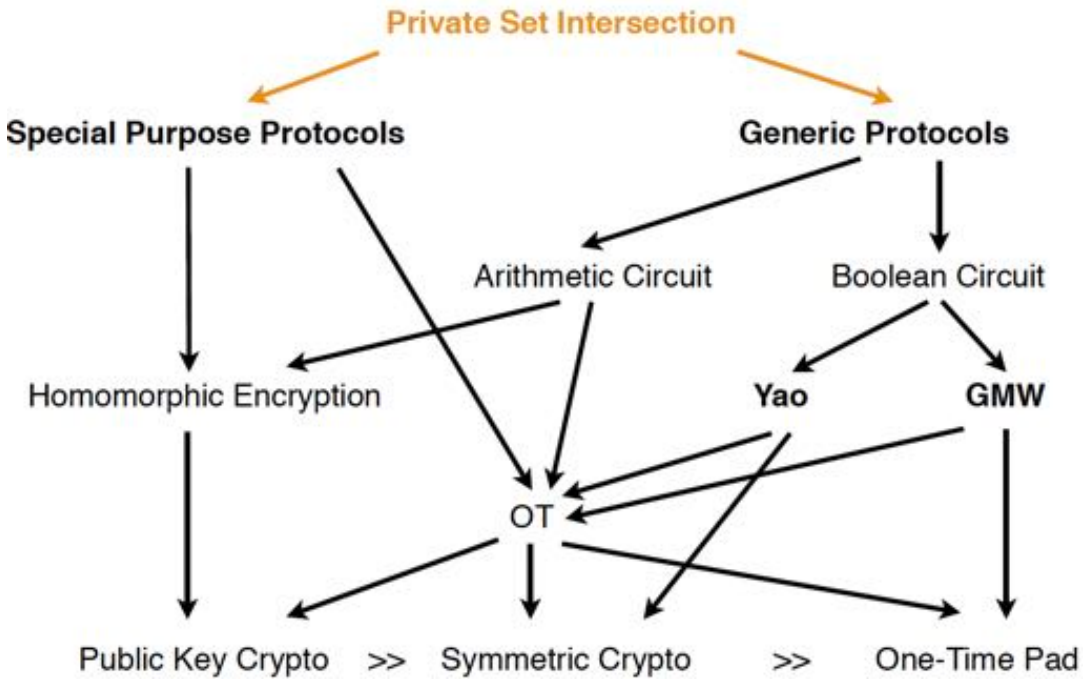

There are generic protocols to validate computations on authenticated data without knowledge of the secret key

DNA MSPN
Gates: 298208 Yao Bytes: 9542656 Depth: 615

DNA PSPN
Gates: 228272 Yao Bytes: 7304704 Depth: 589

NIPS MSPN
Gates: 1001477 Yao Bytes: 32047264 Depth: 970

Homomorphic sum-product network

[Molina, Weinert, Treiber, Schneider, Kersting 2019]

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

The Explorative Automatic Statistician

missing value

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

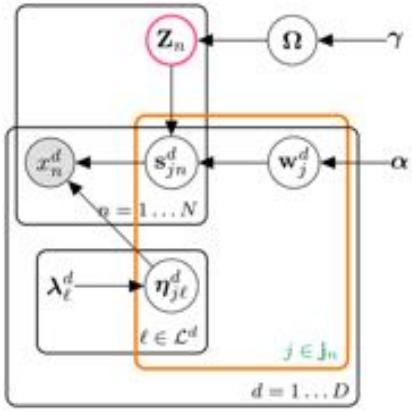

Bayesian Type Discovery

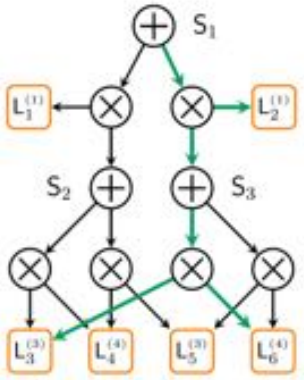

Mixed Sum-Product Network

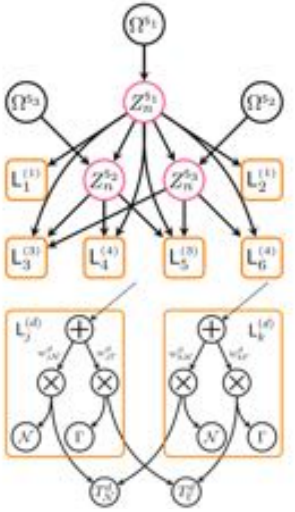

Automatic Statistician

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

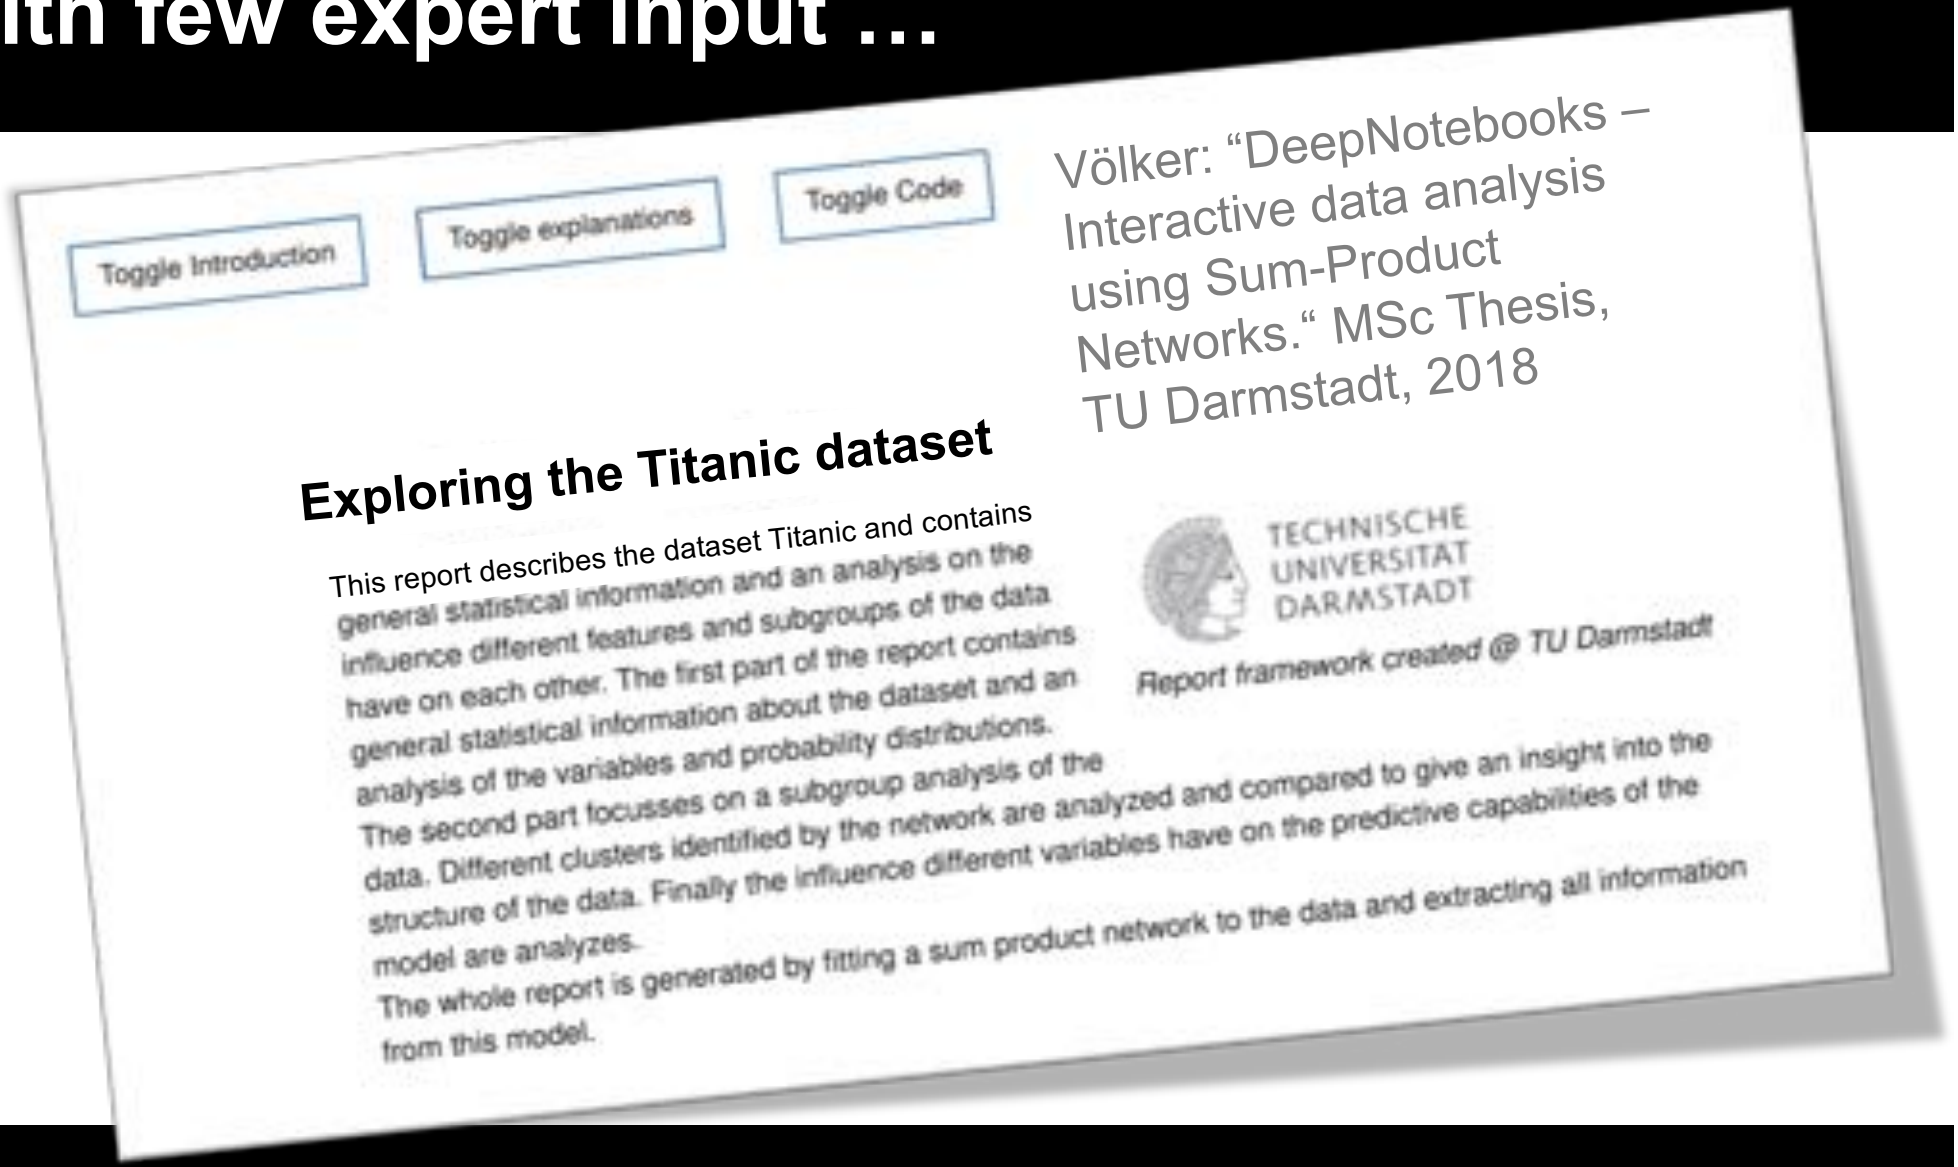

Völker: "DeepNotebooks – Interactive data analysis using Sum-Product Networks." MSc Thesis, TU Darmstadt, 2018

…and can compile data reports automatically

*Similar to [Lapuschkin et al., Nature Communication 10:1096, 2019]

The machine understands the data with no expert input …

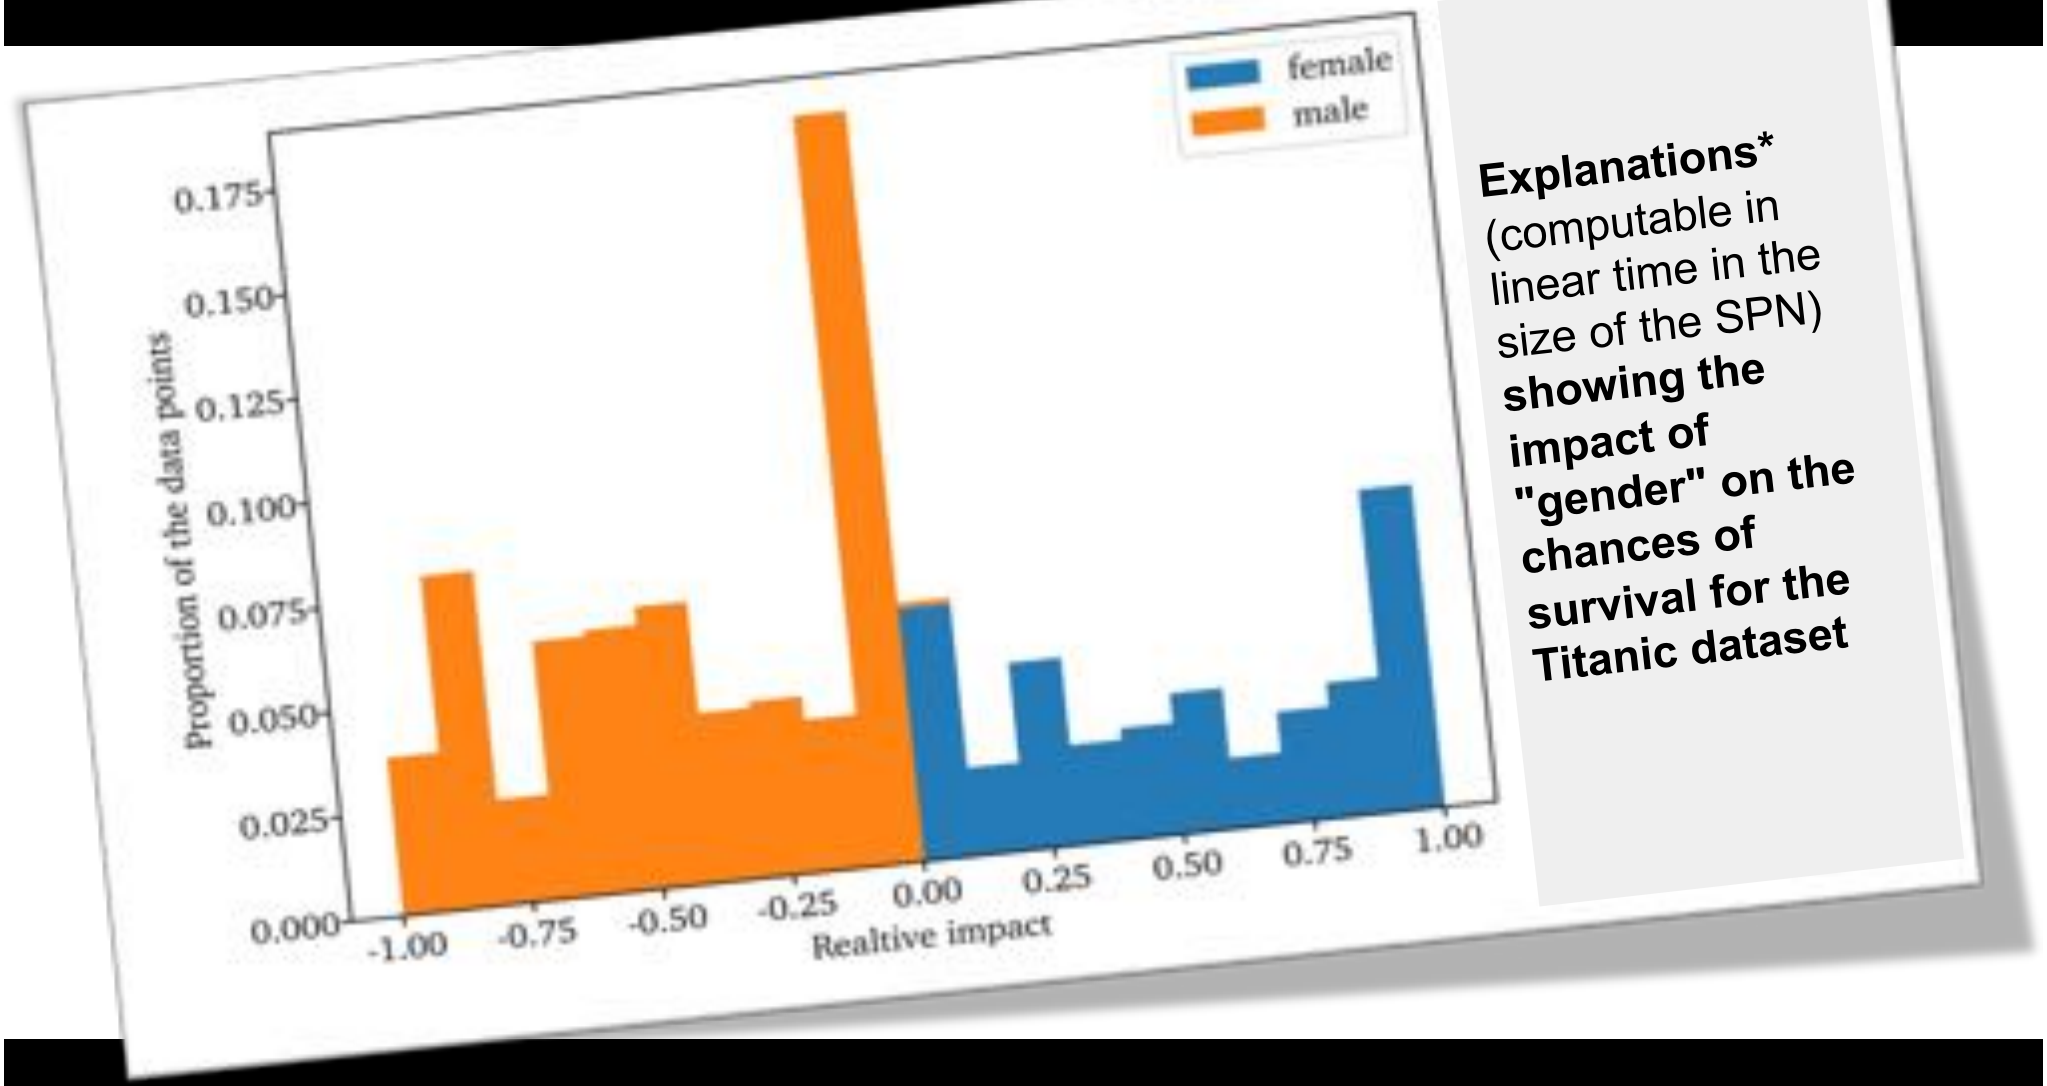

Explanations* (computable in linear time in the size of the SPN) **showing the impact of "gender" on the chances of survival for the Titanic dataset**

…and can compile data reports automatically

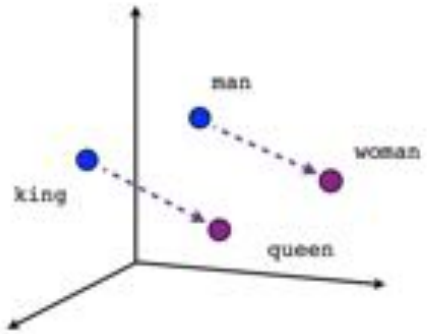

Male-Female

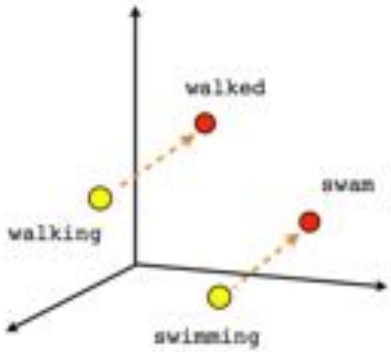

Verb tense

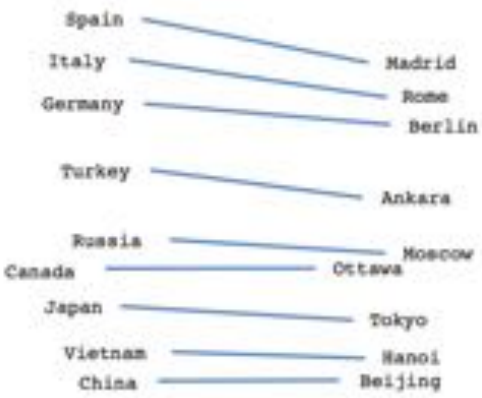

Country-Capital

vector[Queen] = vector[King] - vector[Man] + vector[Woman]

Neural Embeddings

Words and sentences in vector spaces

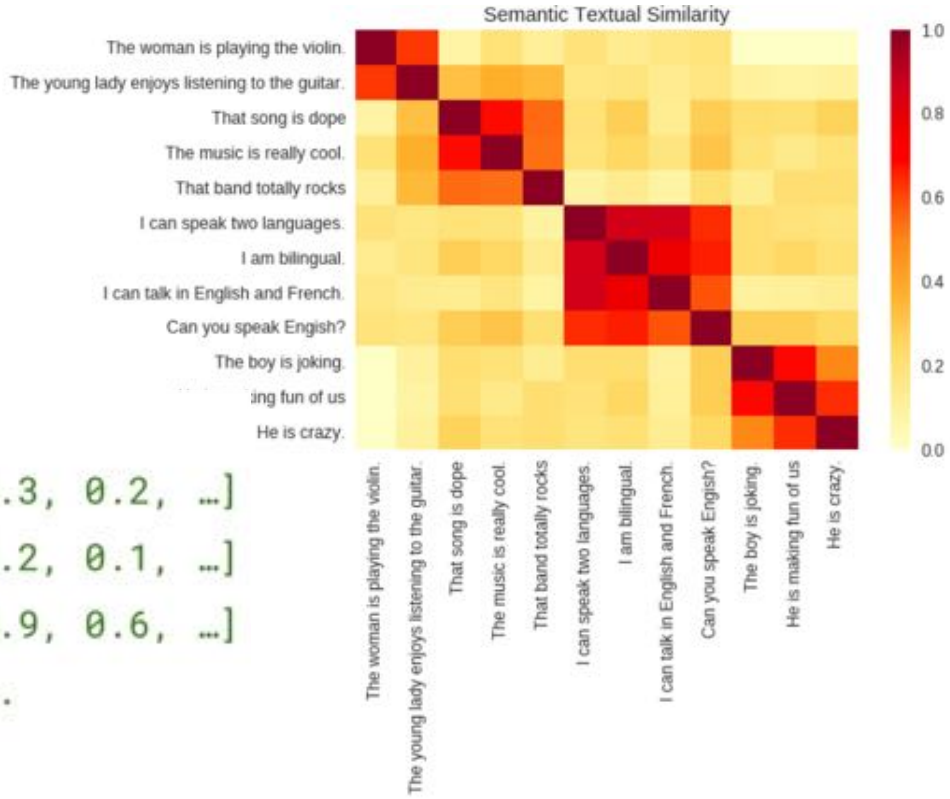

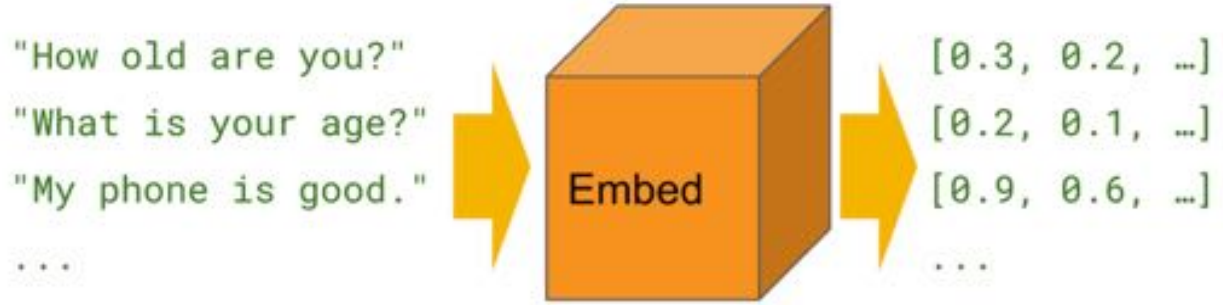

The Moral Choice Machine

Not all stereotypes are bad

[Jentzsch, Schramowski, Rothkopf, Kersting AIES 2019]

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

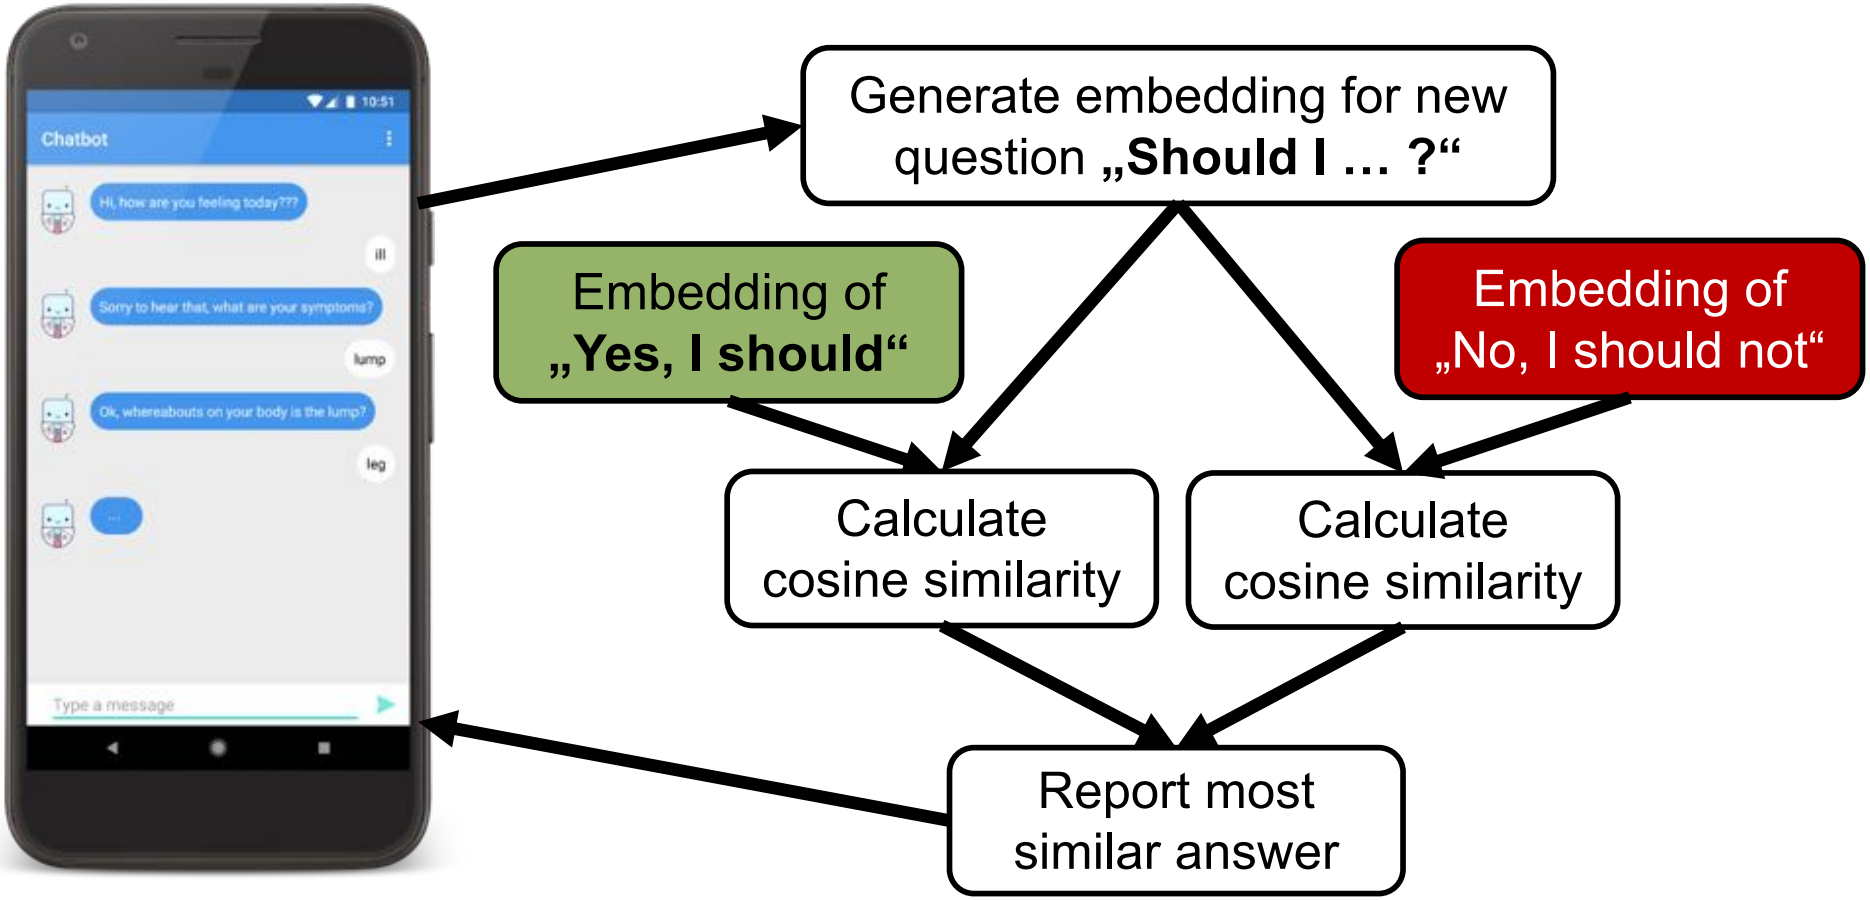

The Moral Choice Machine

Not all stereotypes are bad

[Jentzsch, Schramowski, Rothkopf, Kersting AIES 2019]

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

TECHNISCHE UNIVERSITÄT DARMSTADT

https://www.hr-fernsehen.de/sendungen-a-z/hauptsache-kultur/sendungen/hauptsache-kultur,sendung-56324.html

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.

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 Third Wave of AI

TECHNISCHE UNIVERSITÄT DARMSTADT

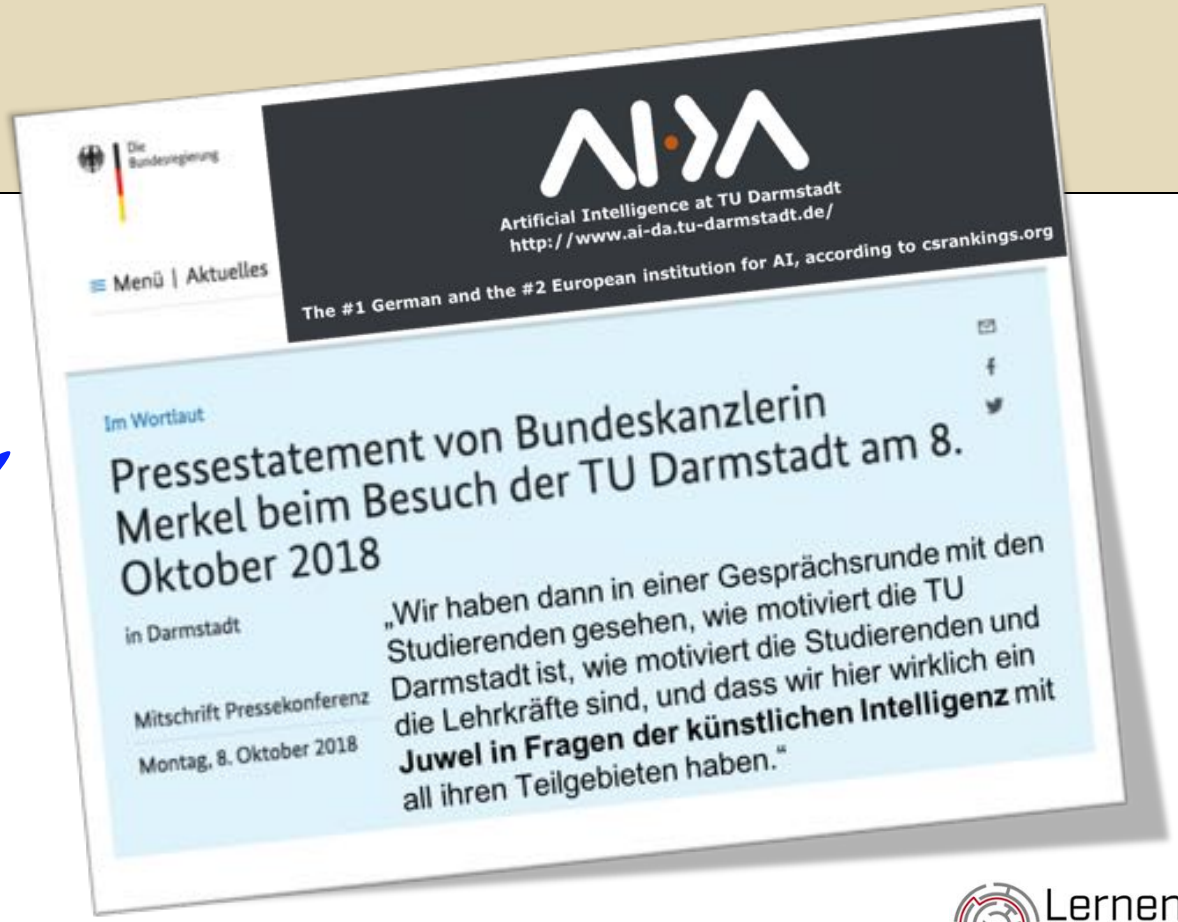

Machine Learning and Artificial Intelligence: Two Fellow Travelers on the Quest for Intelligent Behavior in Machines

Kristian Kersting

Frontiers in Big Data
Published on 19 Nov 2018

OPEN ACCESS

Prof. Dr. Kristian Kersting

Lernende Systeme
DIE PLATTFORM FÜR KÜNSTLICHE INTELLIGENZ

Federal Ministry of Education and Research