Deep machines that know when they do not know

TECHNISCHE UNIVERSITÄT DARMSTADT

Centre for Cognitive Science

Fachbereich Informatik


Die Bundesregierung

≡ Menü | Aktuelles

Artificial Intelligence at TU Darmstadt
http://www.ai-da.tu-darmstadt.de/

The #1 German and the #2 European institution for AI, according to csrankings.org

Im Wortlaut

Pressestatement von Bundeskanzlerin Merkel beim Besuch der TU Darmstadt am 8. Oktober 2018

in Darmstadt

Mitschrift Pressekonferenz

Montag, 8. Oktober 2018

„Wir haben dann in einer Gesprächsrunde mit den Studierenden gesehen, wie motiviert die TU Darmstadt ist, wie motiviert die Studierenden und die Lehrkräfte sind, und dass wir hier wirklich ein **Juwel in Fragen der künstlichen Intelligenz** mit all ihren Teilgebieten haben."

Prof. Dr. Kristian Kersting

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

ELLIS
European Laboratory for Learning and Intelligent Systems

CLAIRE
CONFEDERATION OF LABORATORIES FOR ARTIFICIAL INTELLIGENCE RESEARCH IN EUROPE

Lernende Systeme
DIE PLATTFORM FÜR KÜNSTLICHE INTELLIGENZ

Federal Ministry of Education and Research

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

TECHNISCHE UNIVERSITÄT DARMSTADT

Centre for Cognitive Science

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

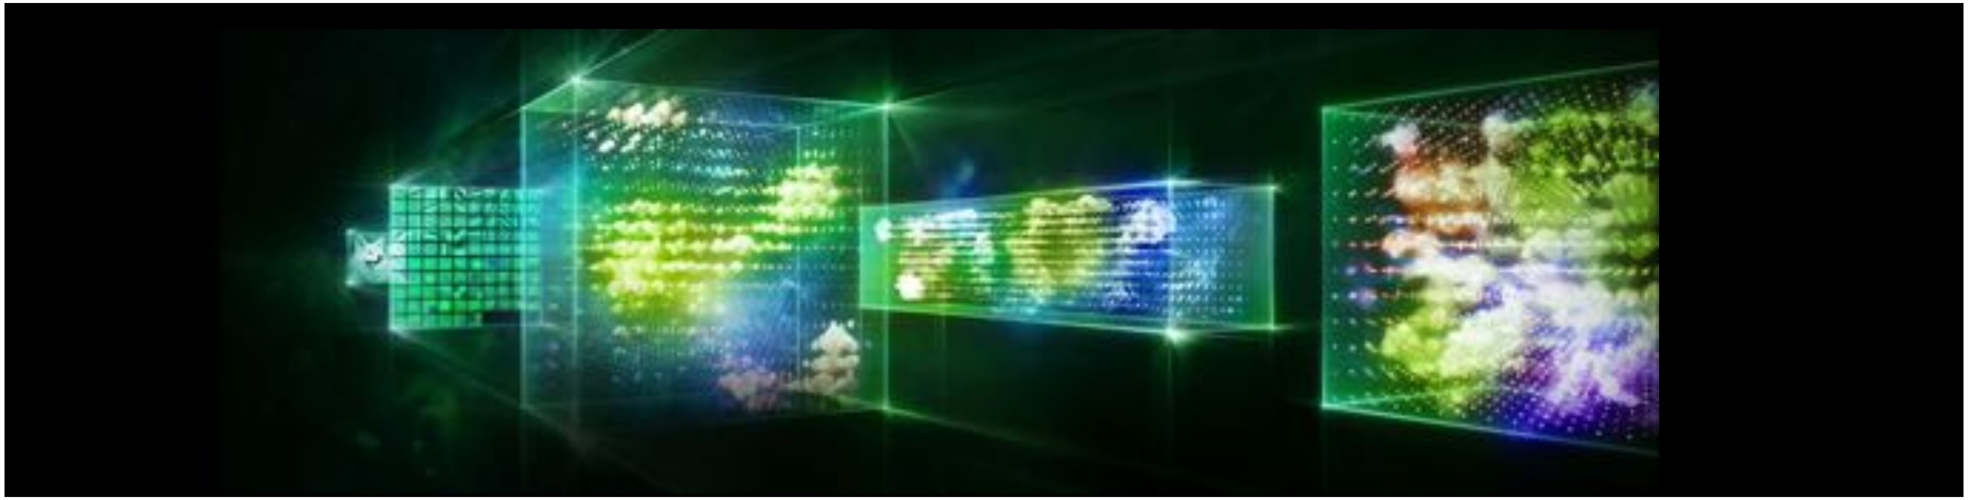

Deep Probabilistic Learning UBER AI Labs

The third wave of AI

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

However, data and learning are only two pieces in the AI puzzle

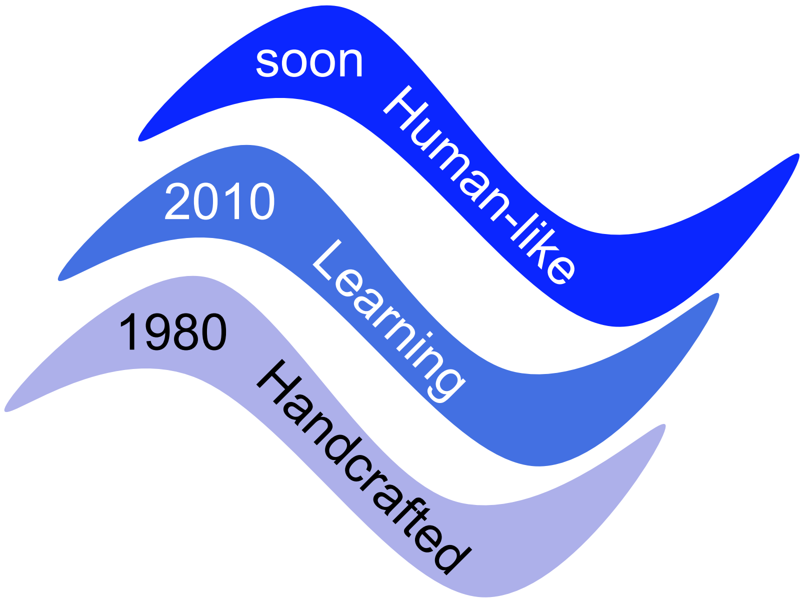

AI systems that can acquire human-like communication and reasoning capabilities, with the ability to recognise new situations and adapt to them.

Deep Neural Networks

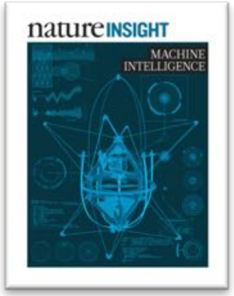

Potentially much more powerful than shallow architectures, represent computations

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

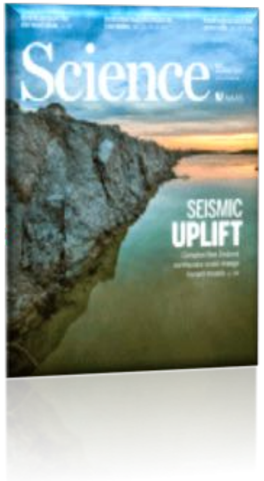


SHARE

1.02k

0

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

They “capture” stereotypes from human language

Video 05:10 Min.

Der Hamster gehört nicht in den Toaster – Wie Forscher von der TU Darmstadt versuchen, Maschinen ... [Videoseite]

hauptsache kultur | 14.03.19, 22:45 Uhr

The Moral Choice Machine

Dos	WEAT	Bias	**Don'ts**	WEAT	Bias
smile	0.116	0.348	rot	-0.099	-1.118
sightsee	0.090	0.281	negative	-0.101	-0.763
cheer	0.094	0.277	harm	-0.110	-0.730
celebrate	0.114	0.264	damage	-0.105	-0.664
picnic	0.093	0.260	slander	-0.108	-0.600
snuggle	0.108	0.238	slur	-0.109	-0.569

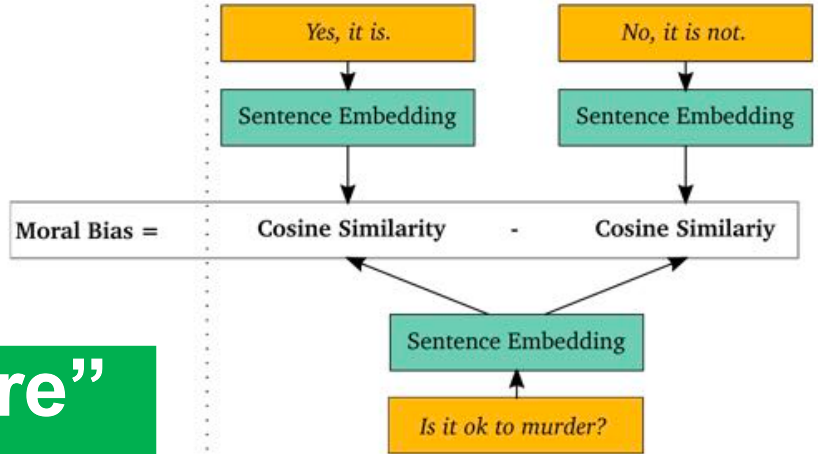

But lucky they also "capture" our moral choices

[Jentzsch, Schramowski, Rothkopf, Kersting AIES 2019]

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

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]

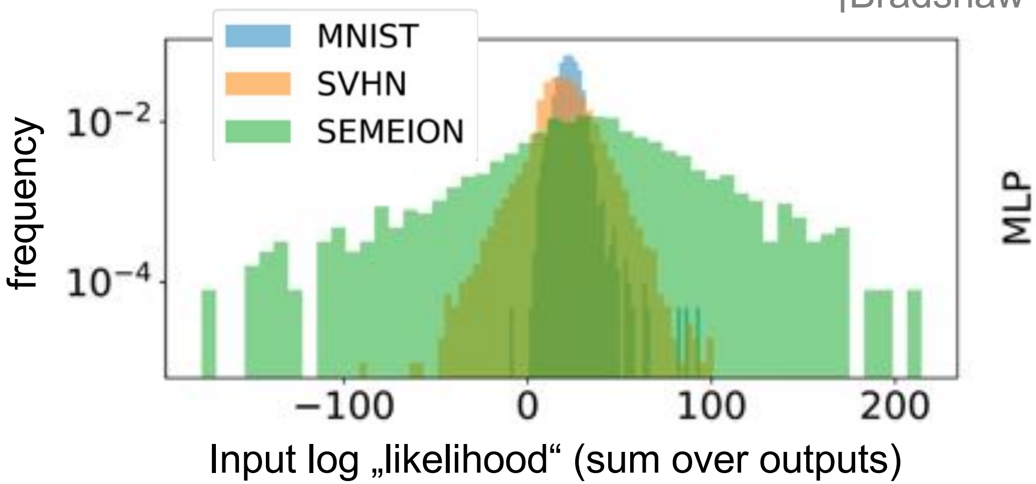

Many DNNs cannot distinguish the datasets

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

The third wave of differentiable programming

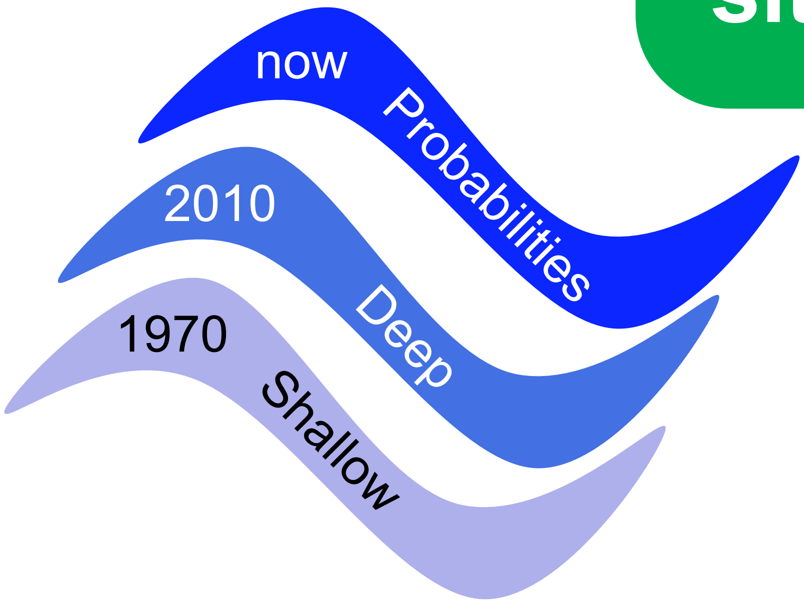

Getting deep systems that know when they do not know and, hence, recognise new situations and adapt to them

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

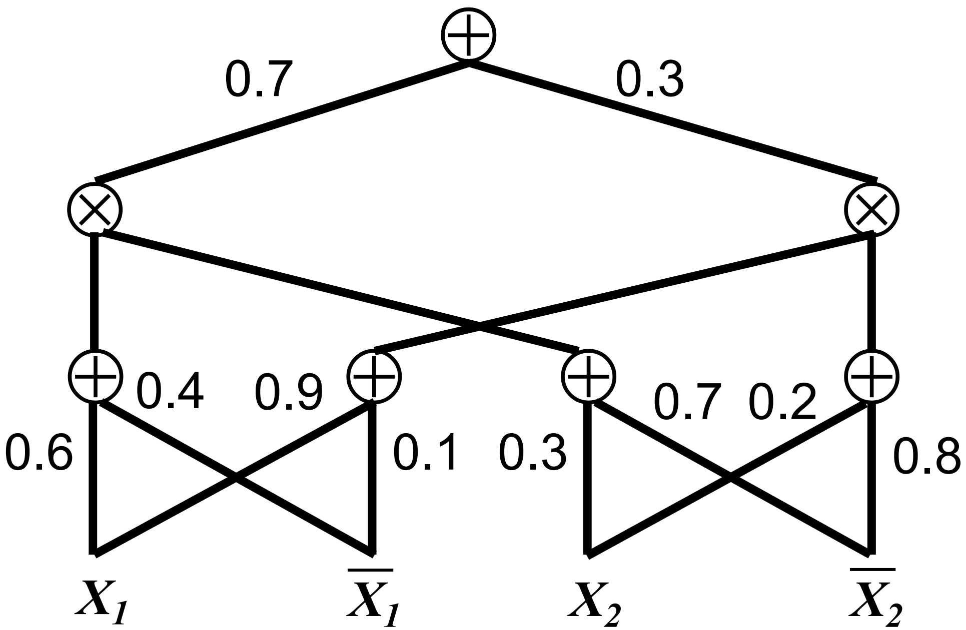

Computational graph (kind of TensorFlow graphs) that encodes how to compute probabilities

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

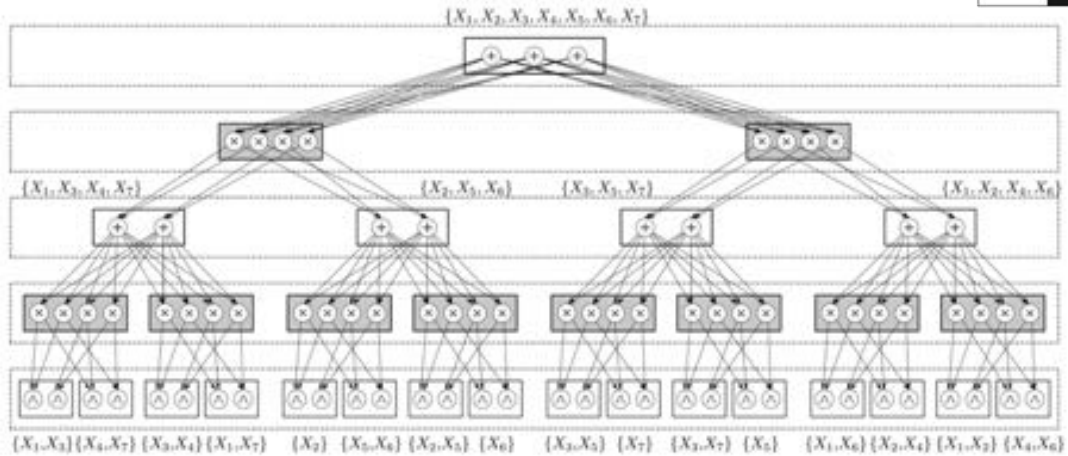

Build a random SPN structure. This can be done in an informed way or completely at random

		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)

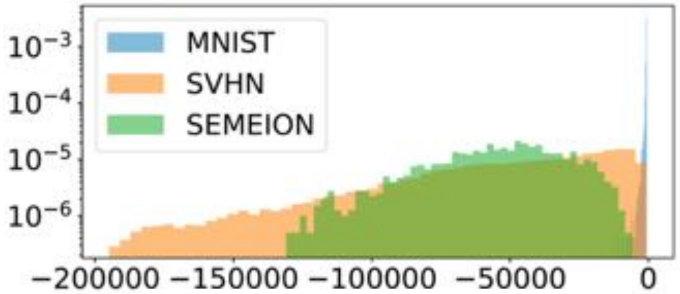

outliers
prototypes
outliers
prototypes

SPNs can have similar predictive performances as (simple) DNNs

SPNs can distinguish the datasets

SPNs know when they do not know by design

[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

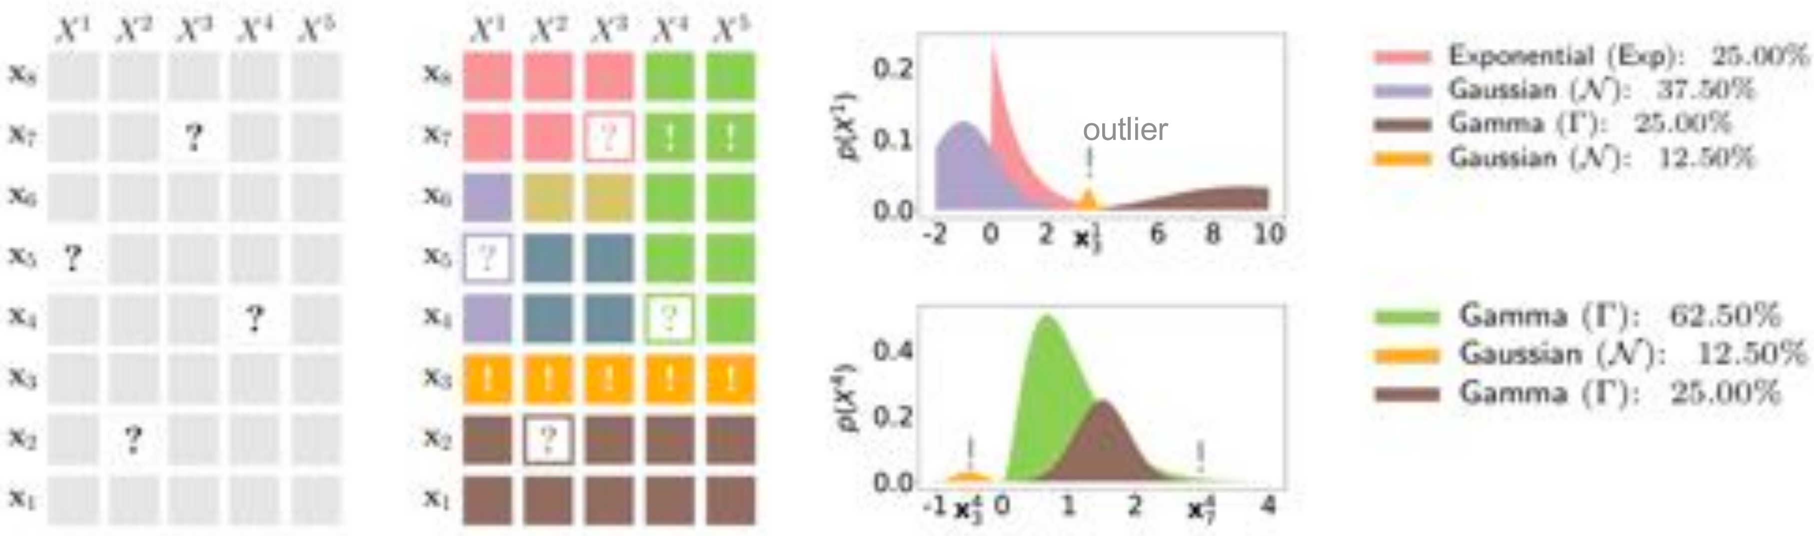

We can even automatically discovers the statistical types and parametric forms of the variables

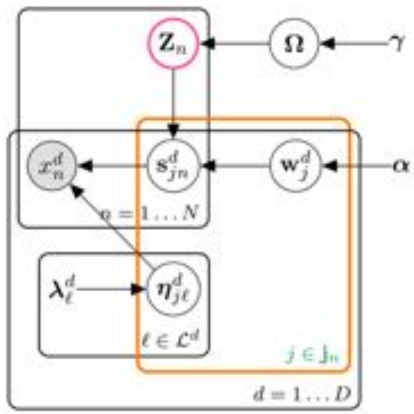

Bayesian Type Discovery

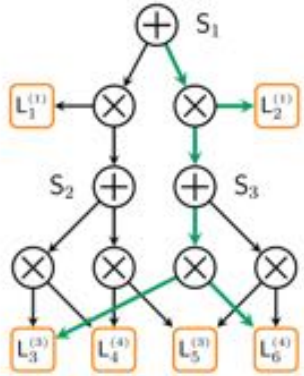

Mixed Sum-Product Network

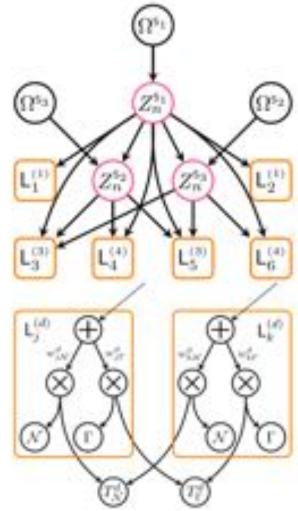

Automatic Statistician

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

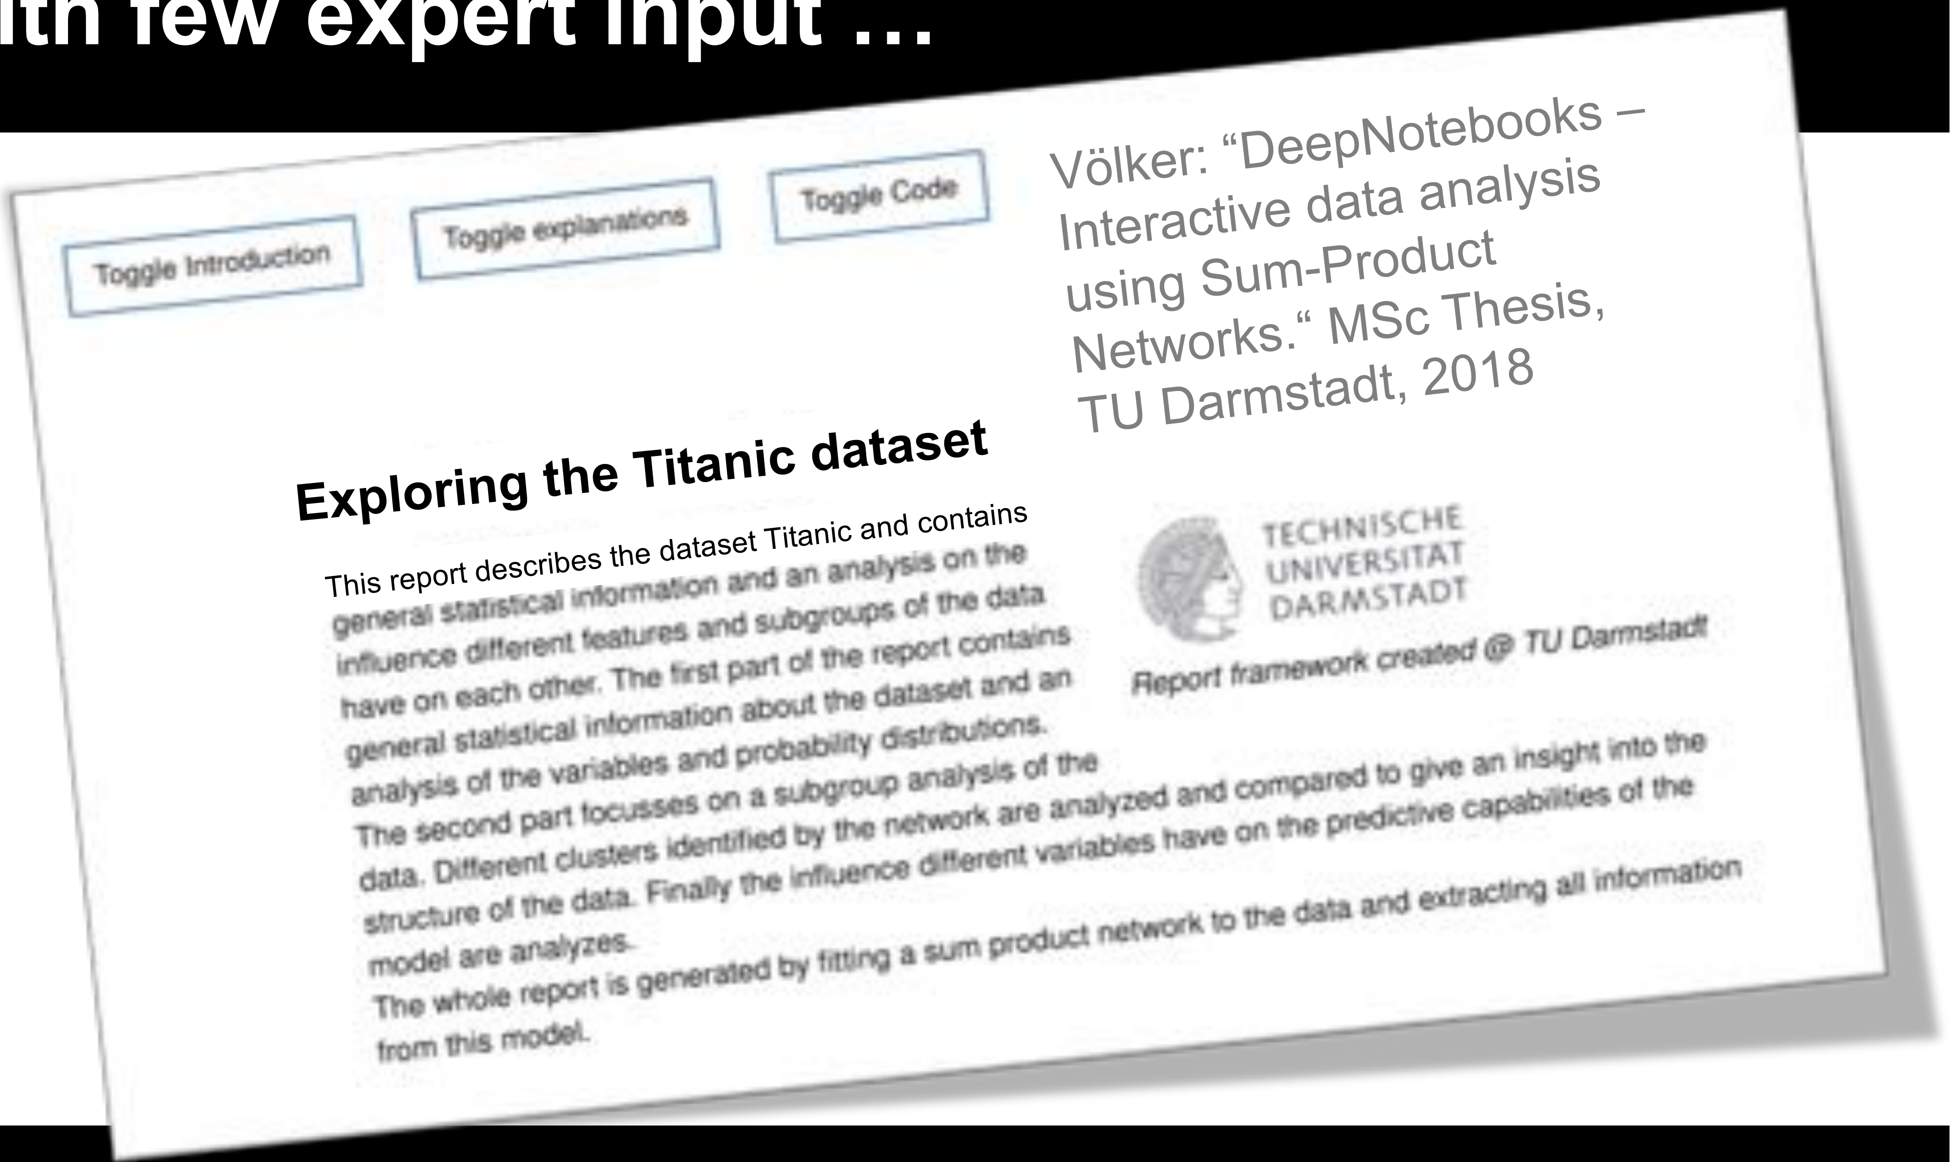
Toggle Introduction Toggle explanations Toggle Code

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.

TECHNISCHE UNIVERSITAT DARMSTADT

Report framework created @ TU Darmstadt

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

…and can compile data reports automatically

Unsupervised scene understanding

[Stelzner, Peharz, Kersting ICML 2019]

UNIVERSITY OF CAMBRIDGE

TECHNISCHE UNIVERSITÄT DARMSTADT

ICML | 2019

Thirty-sixth International Conference on Machine Learning

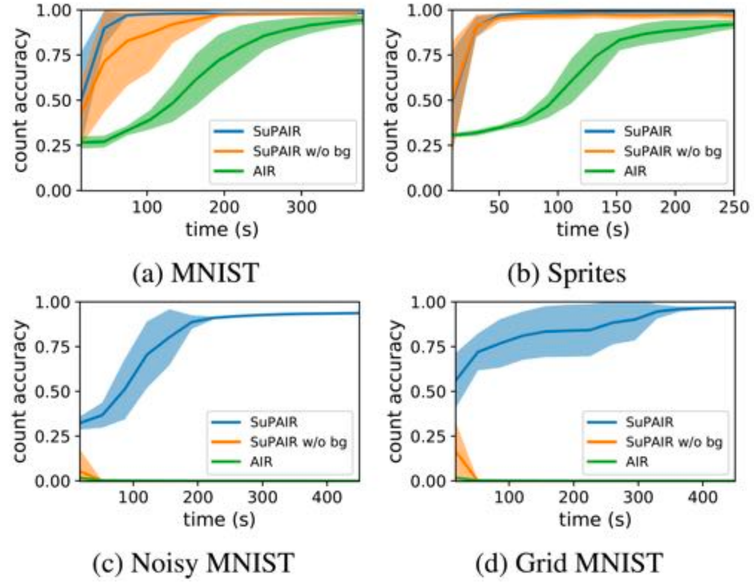

Consider e.g. unsupervised scene understanding using a generative model

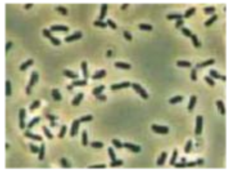

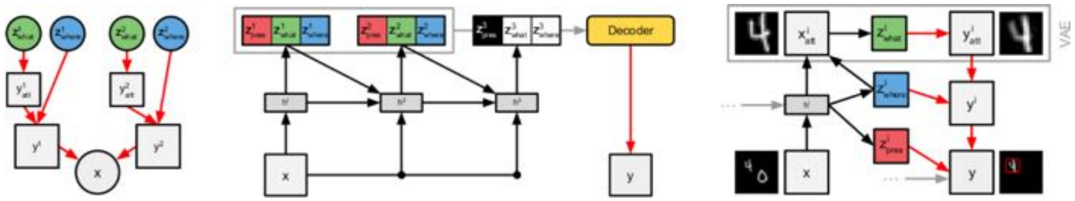

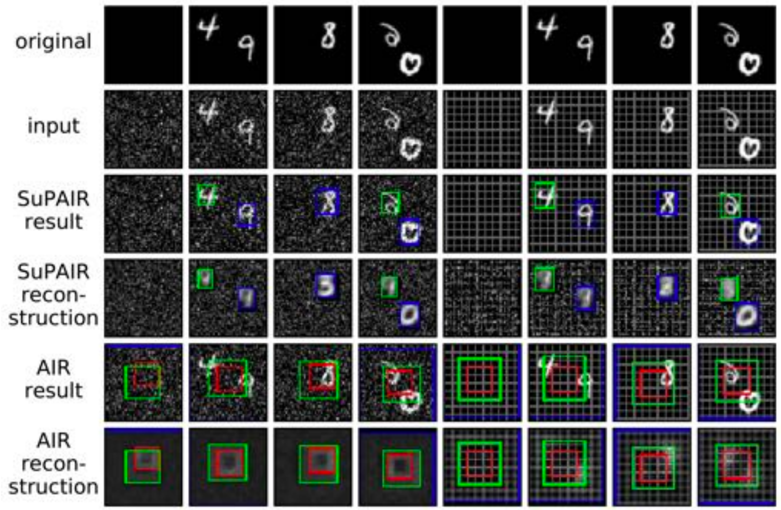

[Attend-Infer-Repeat (AIR) model, Hinton et al. NIPS 2016]

Replace VAE by SPN as object model

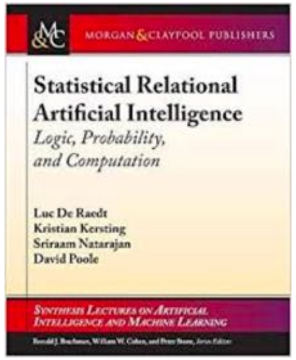

Getting deep systems that reason and know what they don’t know

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“

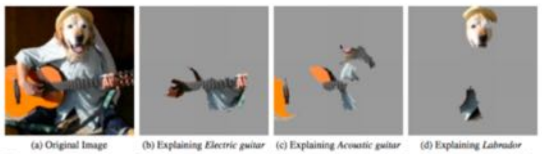

Figure 4: Explaining an image classification prediction made by Google’s Inception network, highlighting positive pixels. The top 3 classes predicted are “Electric Guitar” ($p = 0.32$), “Acoustic guitar” ($p = 0.24$) and “Labrador” ($p = 0.21$)

Teso, Kersting AIES 2019

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

Human algorithms teaches AI a lot

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

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

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

Deep machines that know when they do not know

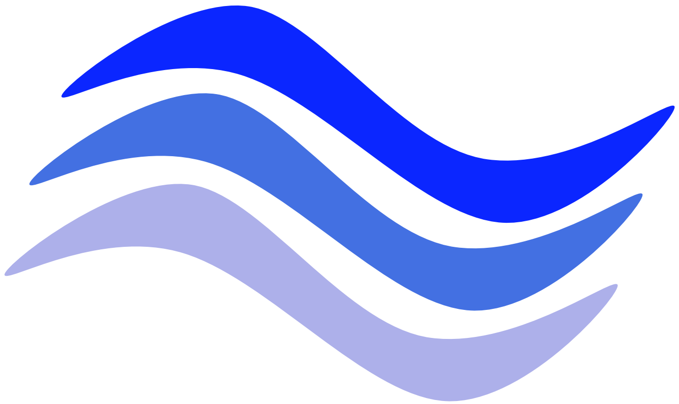

Die Bundesregierung

≡ Menü | Aktuelles

Artificial Intelligence at TU Darmstadt
http://www.ai-da.tu-darmstadt.de/

The #1 German and the #2 European institution for AI, according to csrankings.org

Im Wortlaut

Pressestatement von Bundeskanzlerin Merkel beim Besuch der TU Darmstadt am 8. Oktober 2018

in Darmstadt

Mitschrift Pressekonferenz

Montag, 8. Oktober 2018

„Wir haben dann in einer Gesprächsrunde mit den Studierenden gesehen, wie motiviert die TU Darmstadt ist, wie motiviert die Studierenden und die Lehrkräfte sind, und dass wir hier wirklich ein **Juwel in Fragen der künstlichen Intelligenz** mit all ihren Teilgebieten haben."

**Prof. Dr.
Kristian Kersting**

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

ELLIS
European Laboratory for Learning and Intelligent Systems

CLAIRE
CONFEDERATION OF LABORATORIES FOR ARTIFICIAL INTELLIGENCE RESEARCH IN EUROPE

Lernende Systeme
DIE PLATTFORM FÜR KÜNSTLICHE INTELLIGENZ

Federal Ministry of Education and Research

TECHNISCHE UNIVERSITÄT DARMSTADT

Centre for Cognitive Science

Fachbereich Informatik