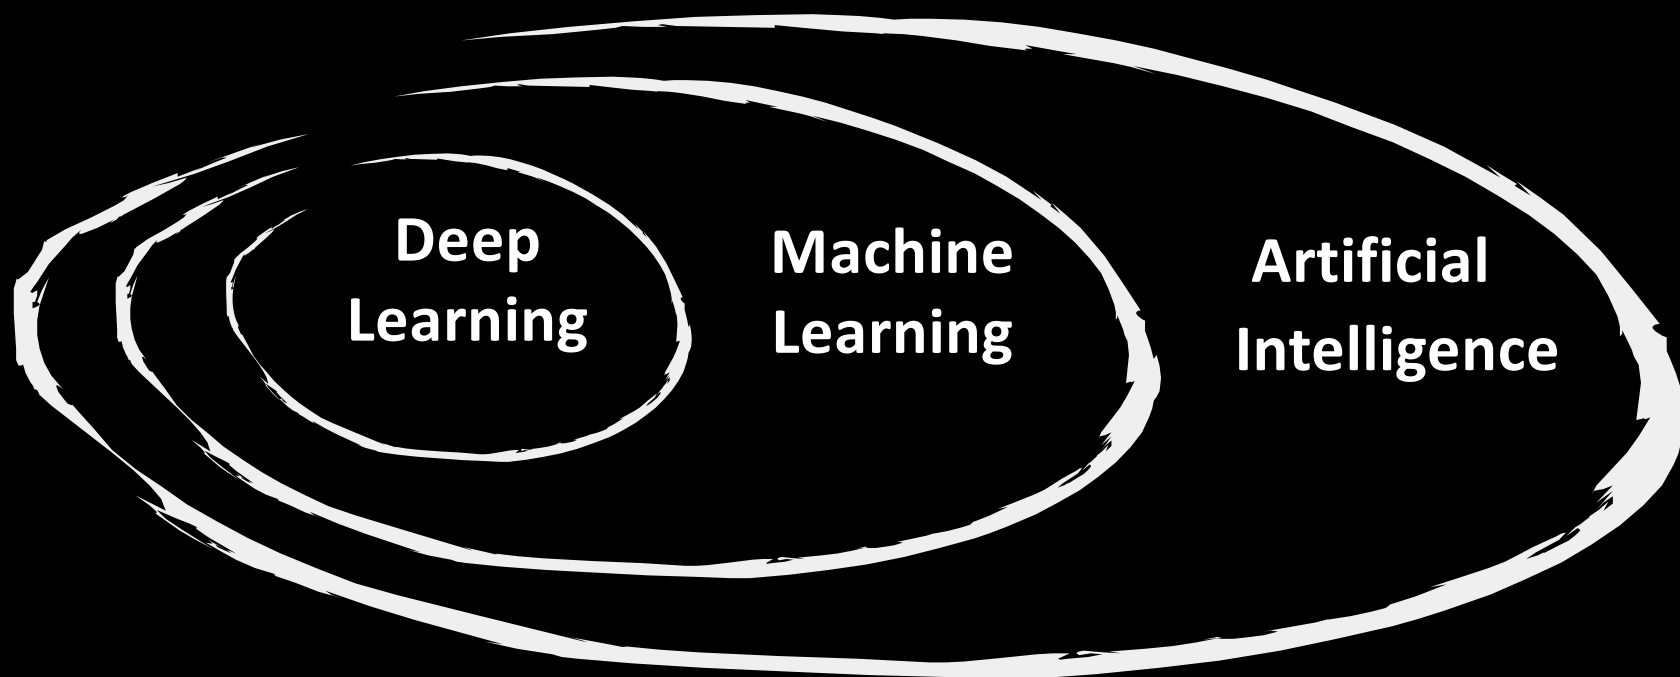


Systems AI

CLAIRE



The computational and mathematical modeling of complex AI systems



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



Kristian Kersting



Deep Neural Networks

Potentially much more powerful than shallow architectures, represent computations

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



SHARE

REPORTS | PSYCHOLOGY



1.02k



0

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

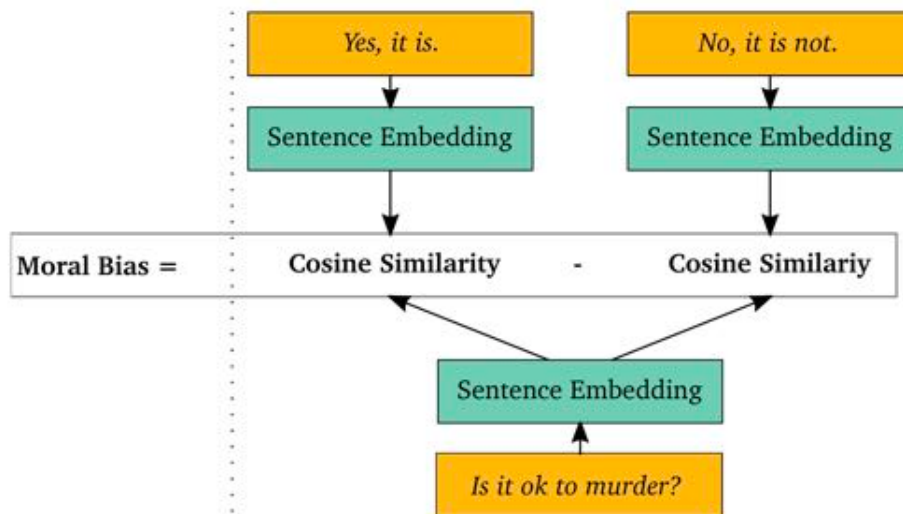
Science 14 Apr 2017:
Vol. 356, Issue 6334, pp. 183-186
DOI: 10.1126/science.aal4230

Deep Neural Networks

Potentially much more powerful than shallow architectures, represent computations

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

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
hug	0.115	0.233	contaminate	-0.102	-0.544
brunch	0.103	0.225	brutalise	-0.118	-0.529
gift	0.130	0.186	poison	-0.131	-0.520
serenade	0.094	0.186	murder	-0.114	-0.515

[Jentzsch, Schramowski, Rothkopf, Kersting 2018]

MNIST



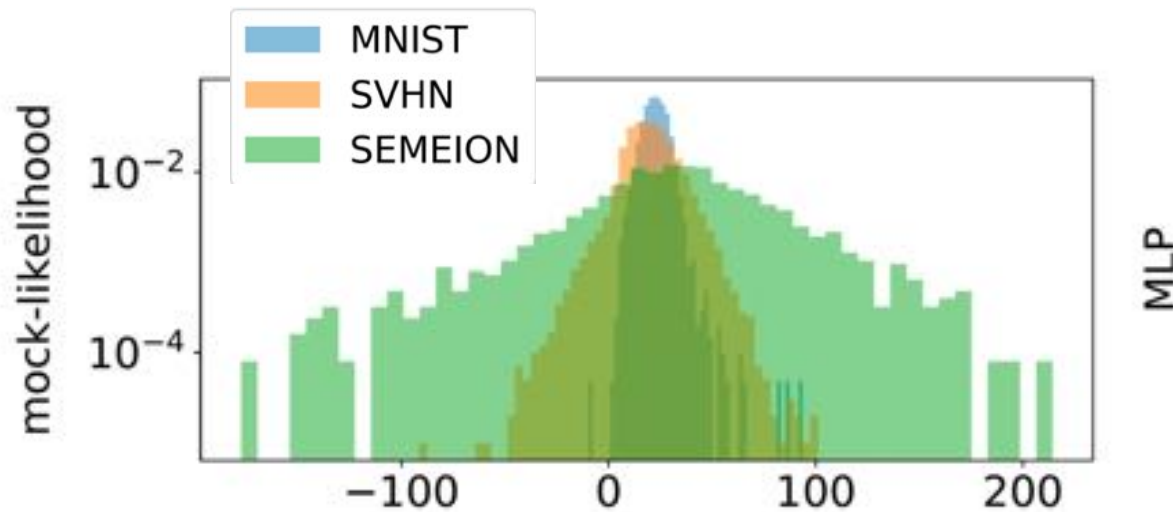
Train & Evaluate

SVHN

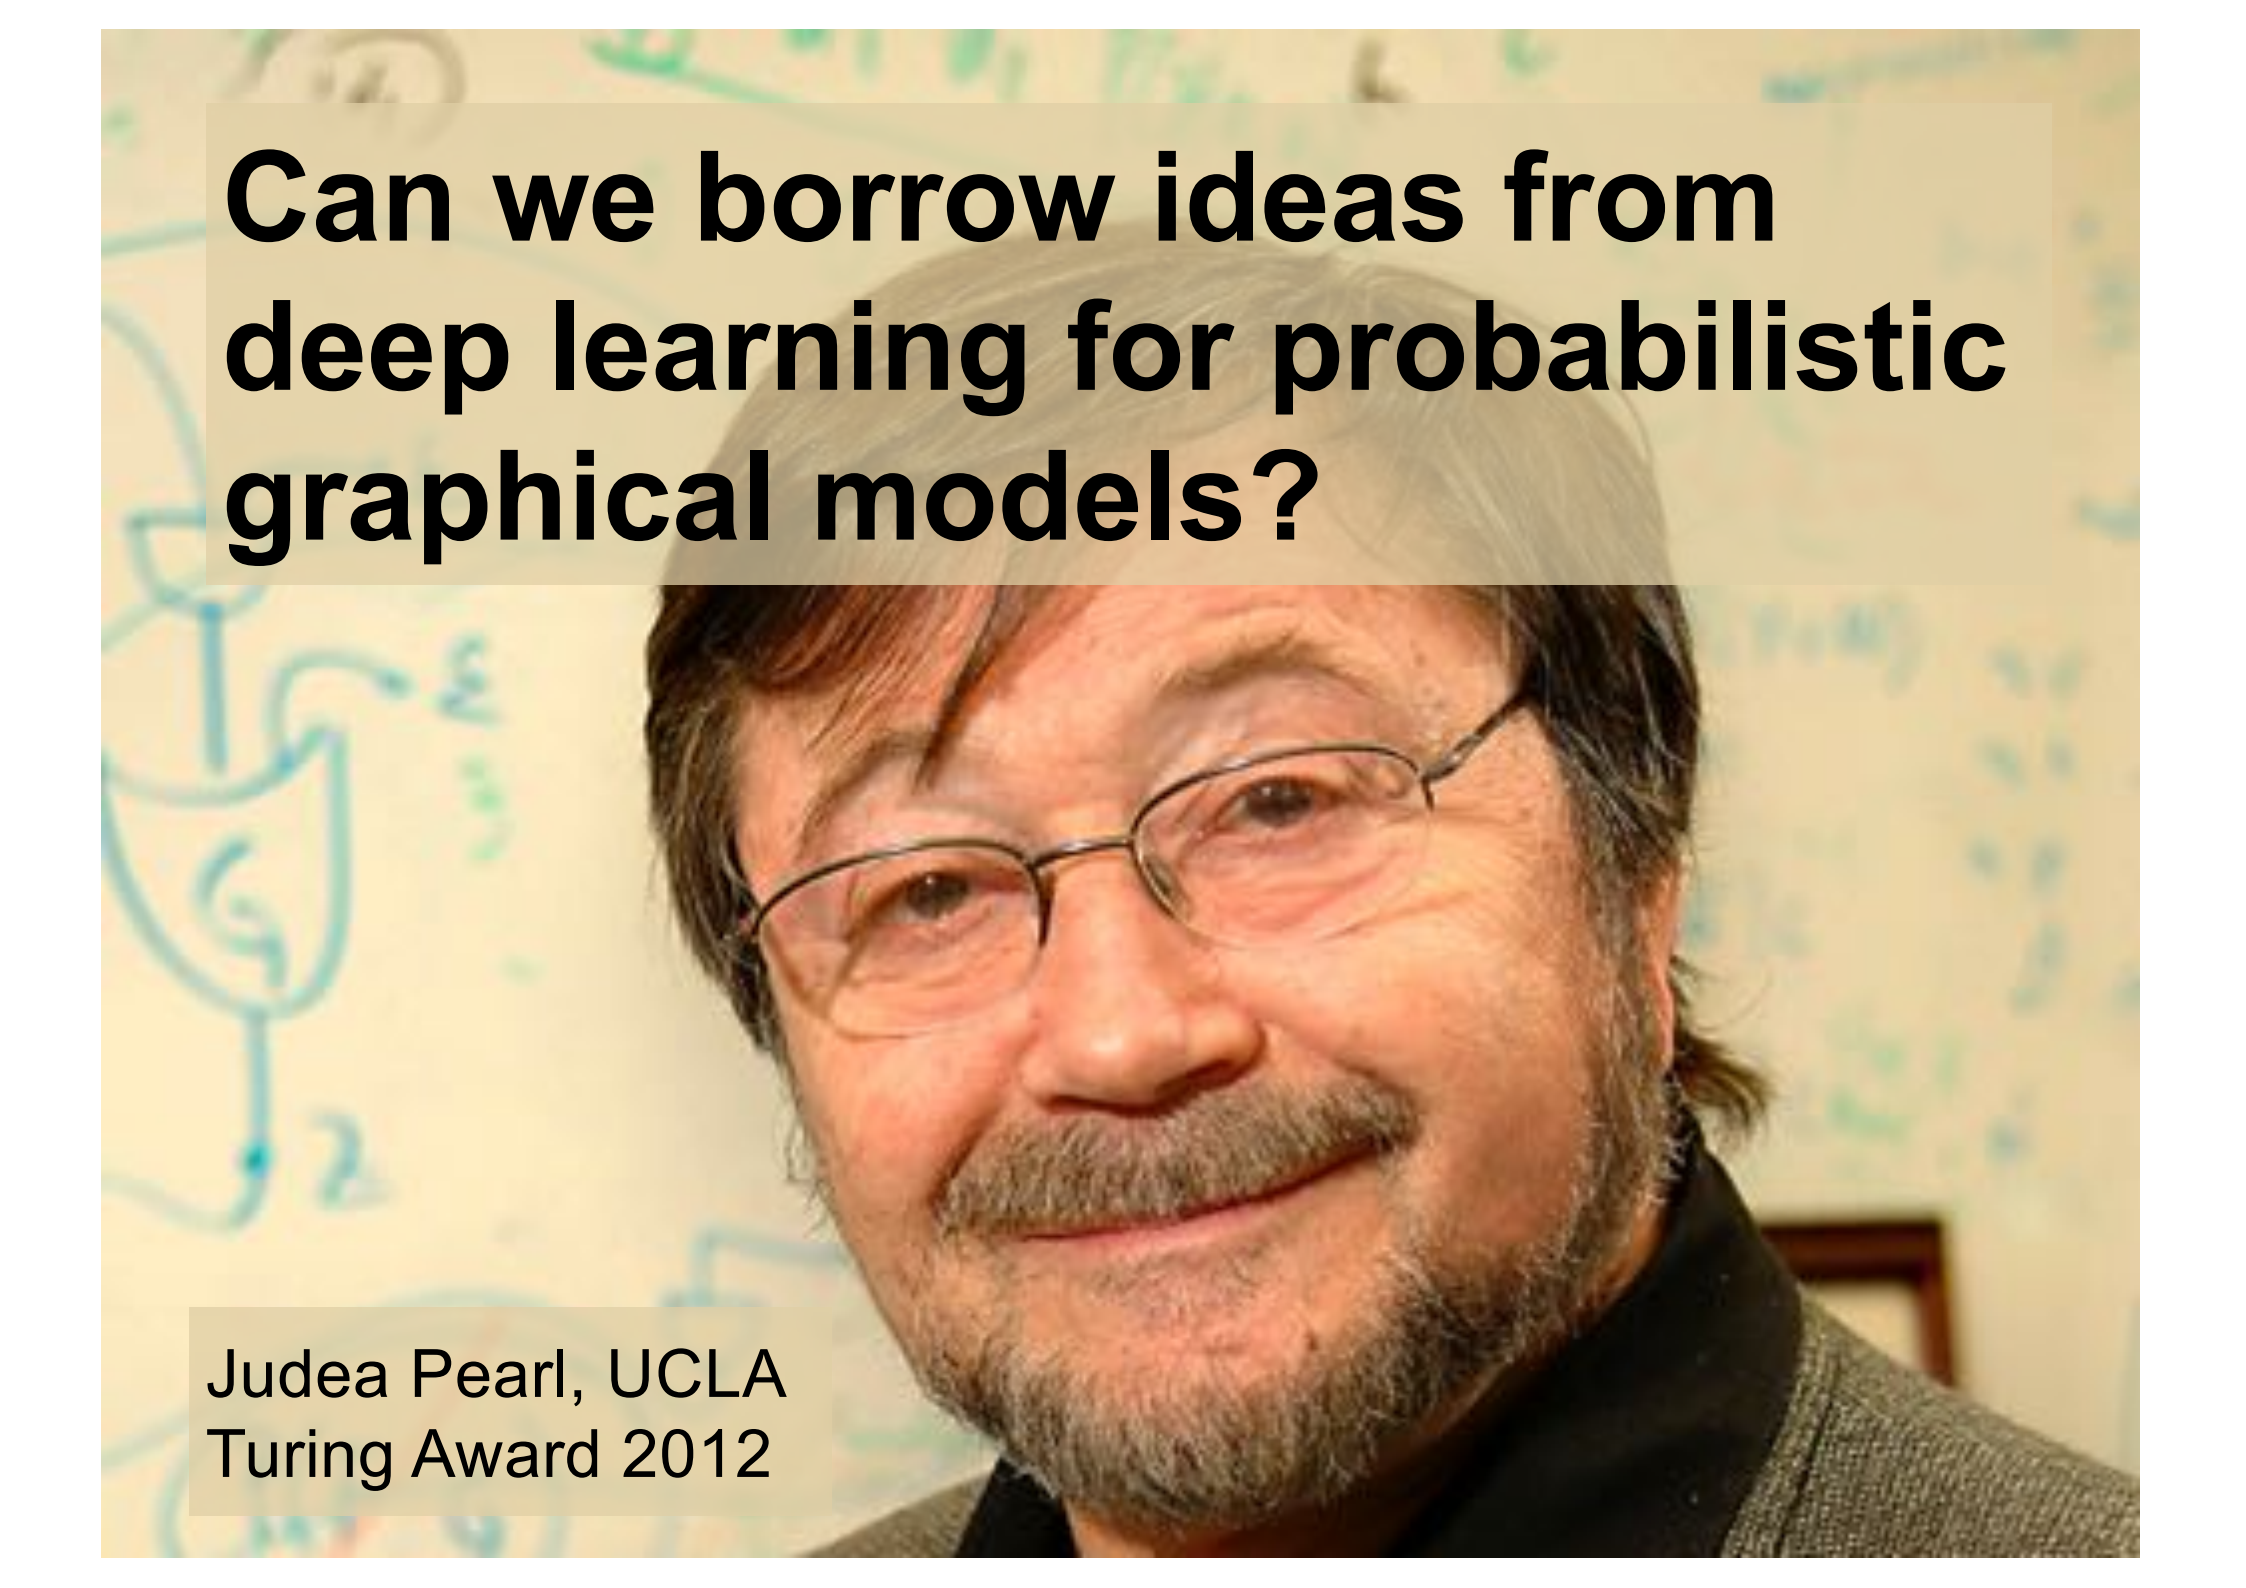


Evaluate

SEMEION



Deep neural networks may not be faithful probabilistic models



Can we borrow ideas from deep learning for probabilistic graphical models?



Judea Pearl, UCLA
Turing Award 2012

Deep Probabilistic Modelling using Sum-Product Networks

Adnan
Darwiche
UCLA



Pedro
Domingos
UW



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

Inference is Linear in Size of Network



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

[Molina, Vergari, Stelzner, Pecharz, Poupard, Di Mauro Kersting 2018]



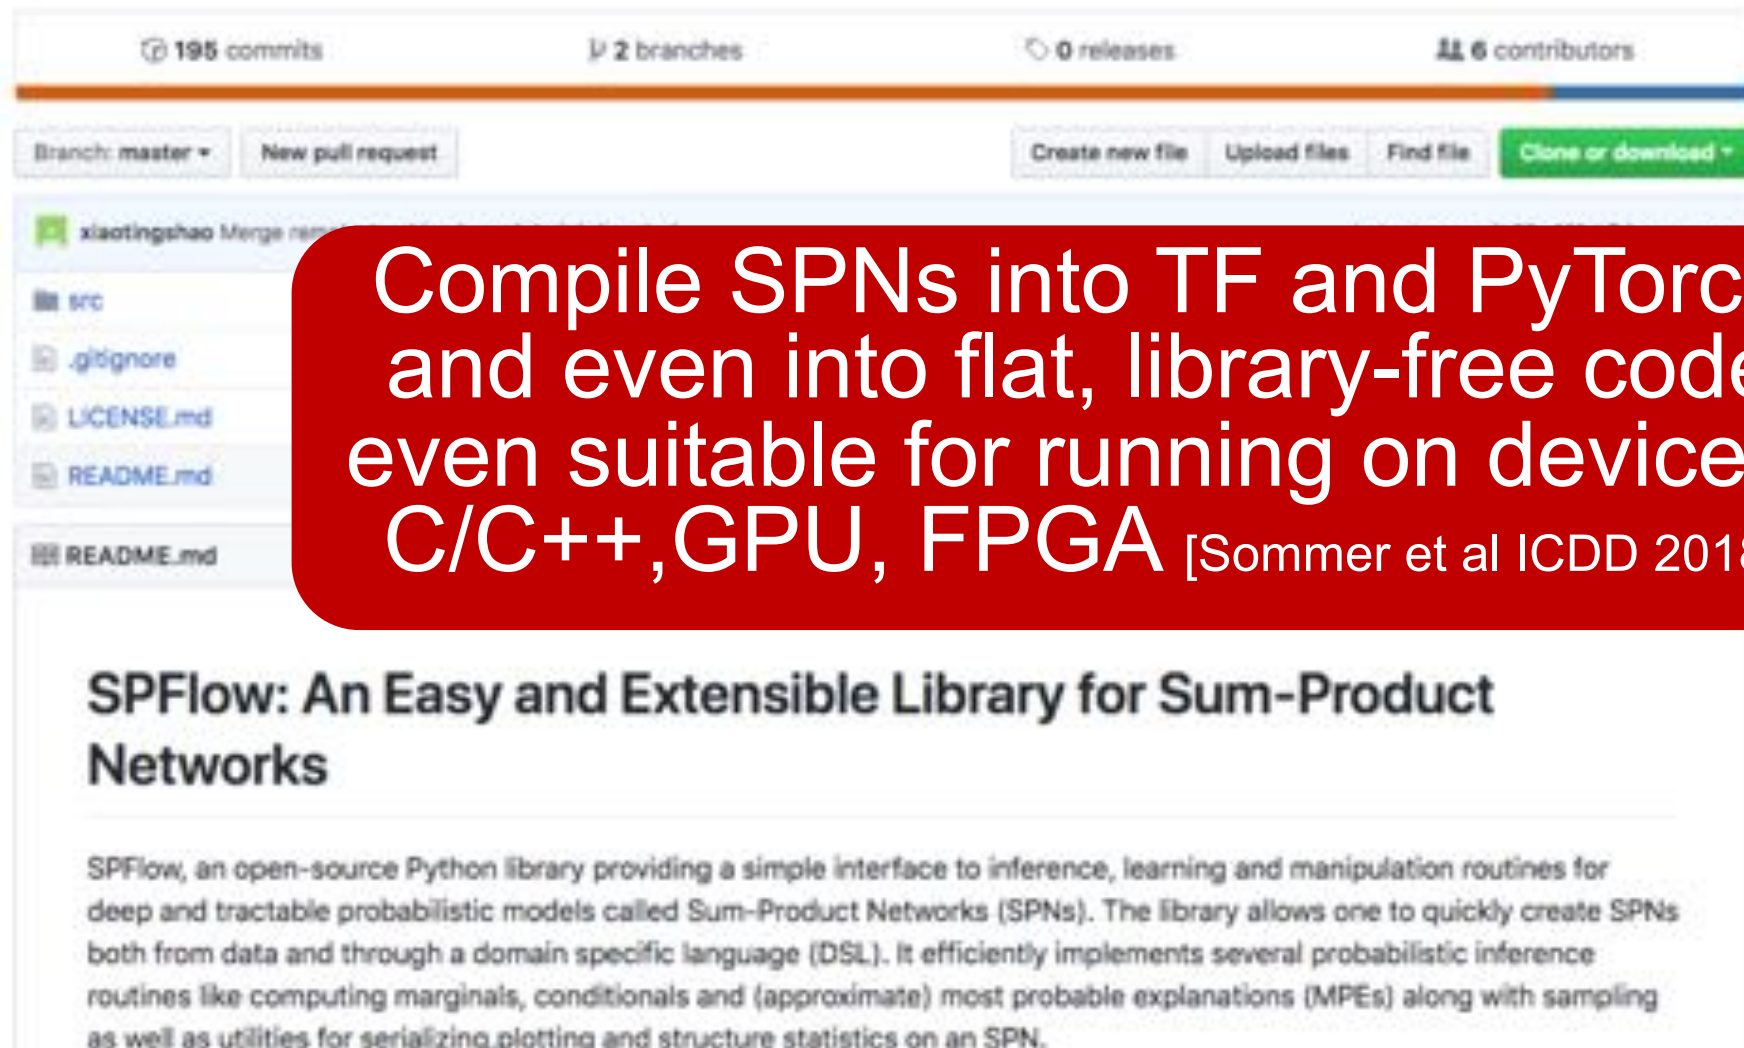
UBER AI Labs



UNIVERSITY OF CAMBRIDGE



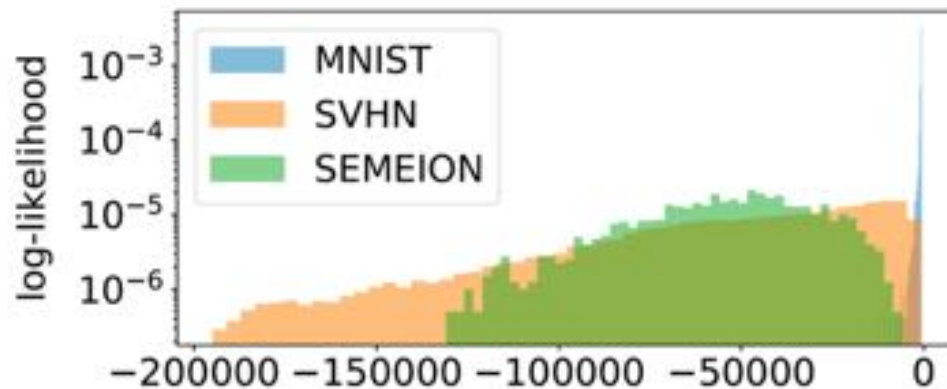

<https://github.com/SPFlow/SPFlow>





Compile SPNs into TF and PyTorch and even into flat, library-free code even suitable for running on devices: C/C++, GPU, FPGA [Sommer et al ICDD 2018]

Random sum-product networks


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



The Explorative Automatic Statistician




- █ Exponential (Exp): 25.00%
- █ Gaussian (\mathcal{N}): 37.50%
- █ Gamma (Γ): 25.00%
- █ Gaussian (\mathcal{N}): 12.50%




- █ Gamma (Γ): 62.50%
- █ Gaussian (\mathcal{N}): 12.50%
- █ Gamma (Γ): 25.00%


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



(a) Graphical model



(b) SPN



(c) Type-augmented SPN

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

The screenshot shows a Jupyter Notebook report interface. At the top, there are three toggle buttons: "Toggle Introduction", "Toggle explanations", and "Toggle Code". The main title of the report is "Exploring the Titanic dataset". The text of the report describes the Titanic dataset and contains general statistical information and an analysis on the influence of different features and subgroups. It mentions that the first part contains general statistical information about the dataset and an analysis of variables and probability distributions. The second part focuses on a subgroup analysis of the data, identifying clusters and analyzing their influence on predictive capabilities. The report is generated by fitting a sum product network to the data and extracting all information from this model. On the right side, there is a citation for Völker's MSc Thesis from TU Darmstadt (2018) and the logo of Technische Universität Darmstadt, with a note that the report framework was created at TU Darmstadt.


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.

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

 TECHNISCHE UNIVERSITÄT DARMSTADT
Report framework created @ TU Darmstadt

...and can compile data reports automatically


P(heart attack | EHR)?



P(heart attack | )?

Statistical Relational AI: Logic, Probability, and Computation (or Bibel meets Bayes)

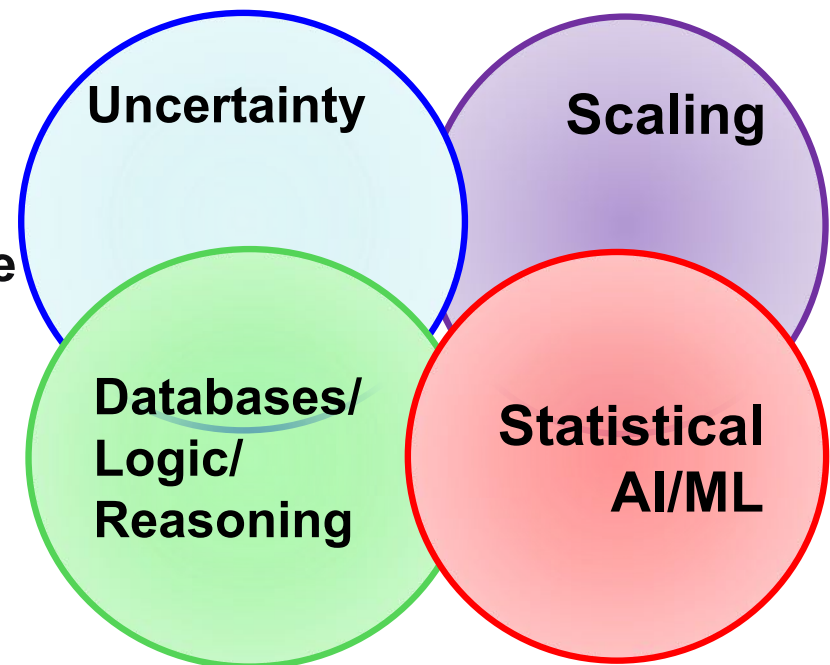
De Raedt, Kersting, Natarajan, Poole: Statistical Relational Artificial Intelligence: Logic, Probability, and Computation. Morgan and Claypool Publishers, ISBN: 9781627058414, 2016.




building general-purpose thinking and learning machines

make the AI/ML expert more effective

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





Probabilistic Models of EHRs


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



TECHNISCHE UNIVERSITÄT DARMSTADT



THE UNIVERSITY OF TEXAS AT DALLAS



Plaque in the left coronary artery

[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

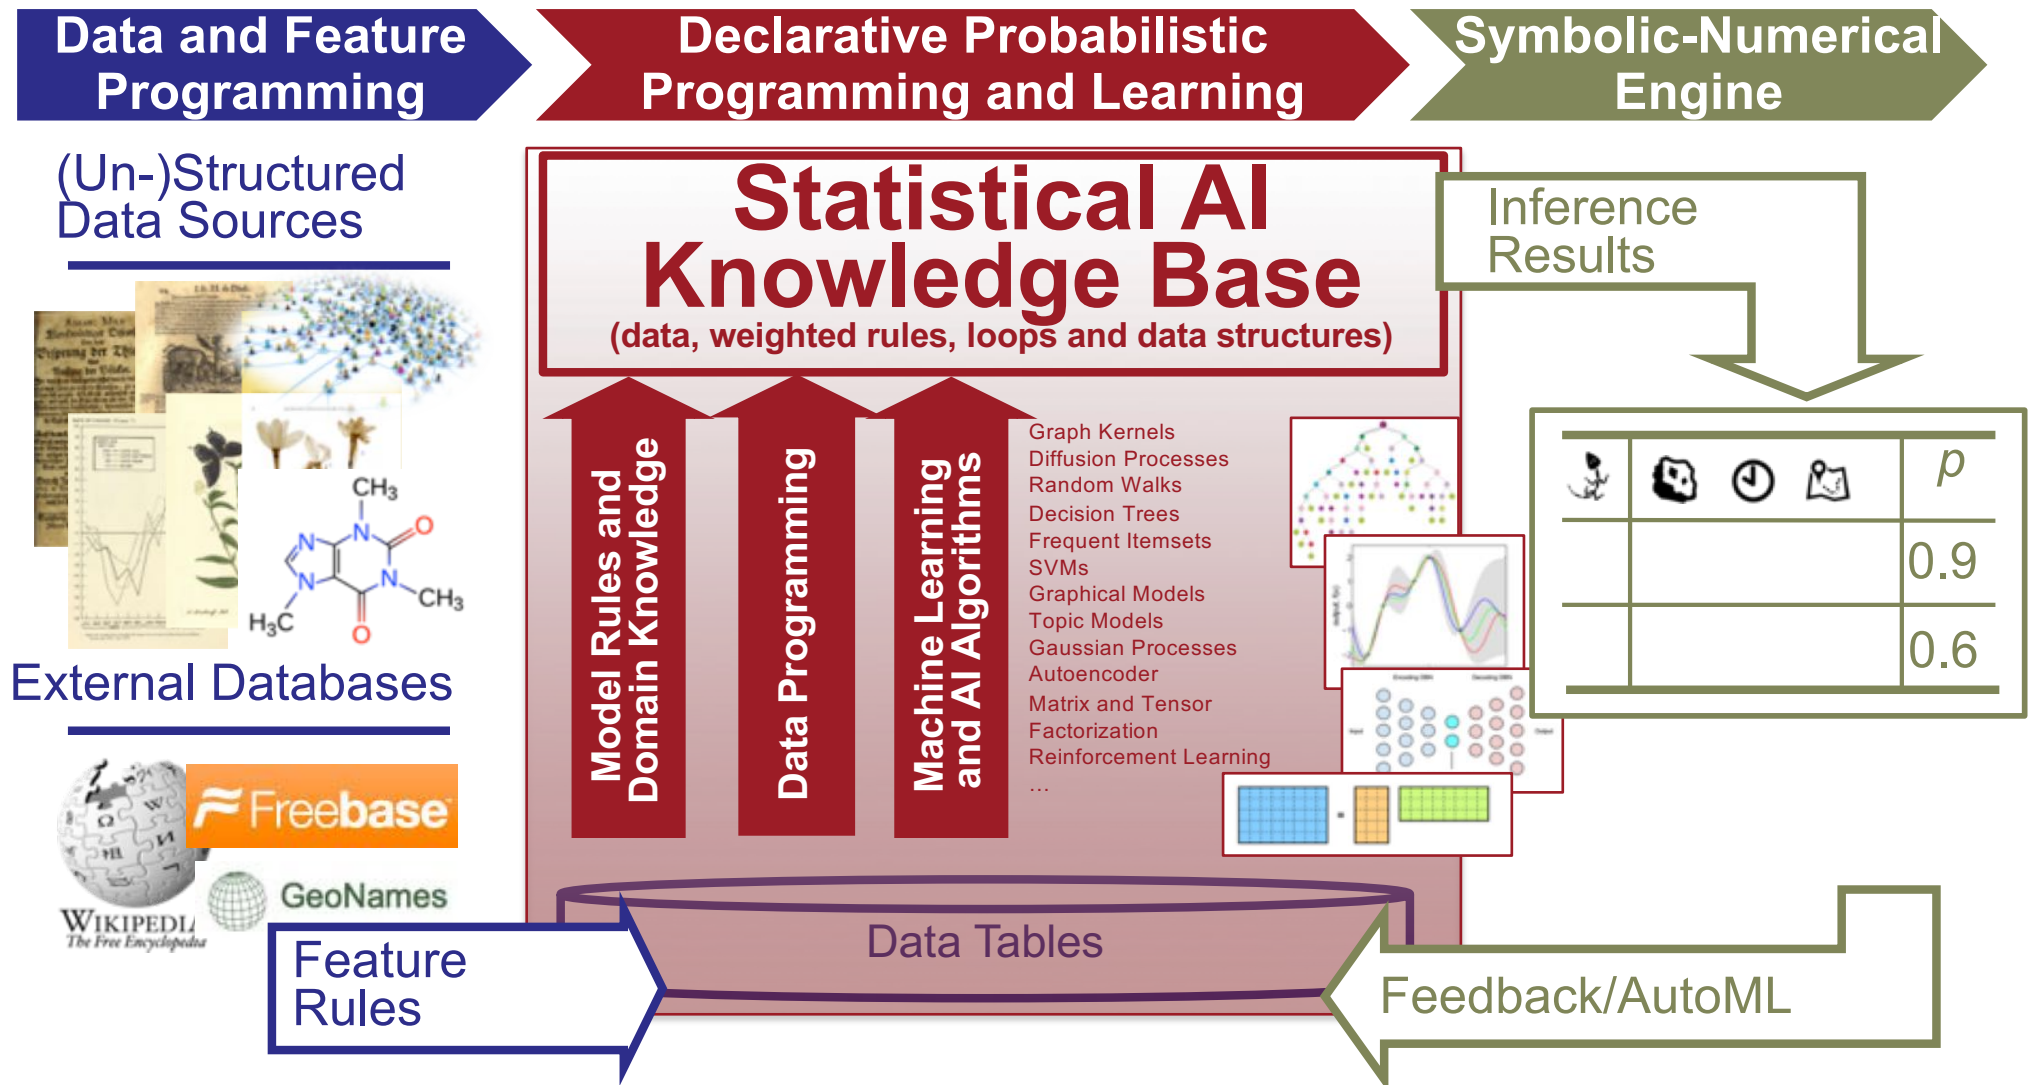
The higher, the better

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] 11%	0.96] 78%	0.93] 50%	9s] 37200x faster
LSM	0.73	0.54	0.62	93 hrs

[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]

This establishes a novel “Deep AI”



And connects well to other communities



Jim Gray Turing Award 1998
“Automated Programming”



Mike Stonebraker Turing Award 2014
“One size does not fit all”

... also Cognitive Science, the twin science of Artificial Intelligence

"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.



Lake, Salakhutdinov, Tenenbaum, Science 350 (6266), 1332-1338, 2015
Tenenbaum, Kemp, Griffiths, Goodman, Science 331 (6022), 1279-1285, 2011

Since we need Systems AI, the computational and mathematical modeling of complex AI systems.

The next breakthrough in AI may not just be a new AI algorithm...

...but may be in the ability to rapidly combine, deploy, and maintain existing algorithms

Since we need Systems AI, the computational and mathematical modeling of complex AI systems.



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>. But also see e.g. **Kordjamshidi, Roth, Kersting**: "Systems AI: A Declarative Learning Based Programming Perspective." IJCAI-ECAI 2018.


There are strong invests into probabilistic programming

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



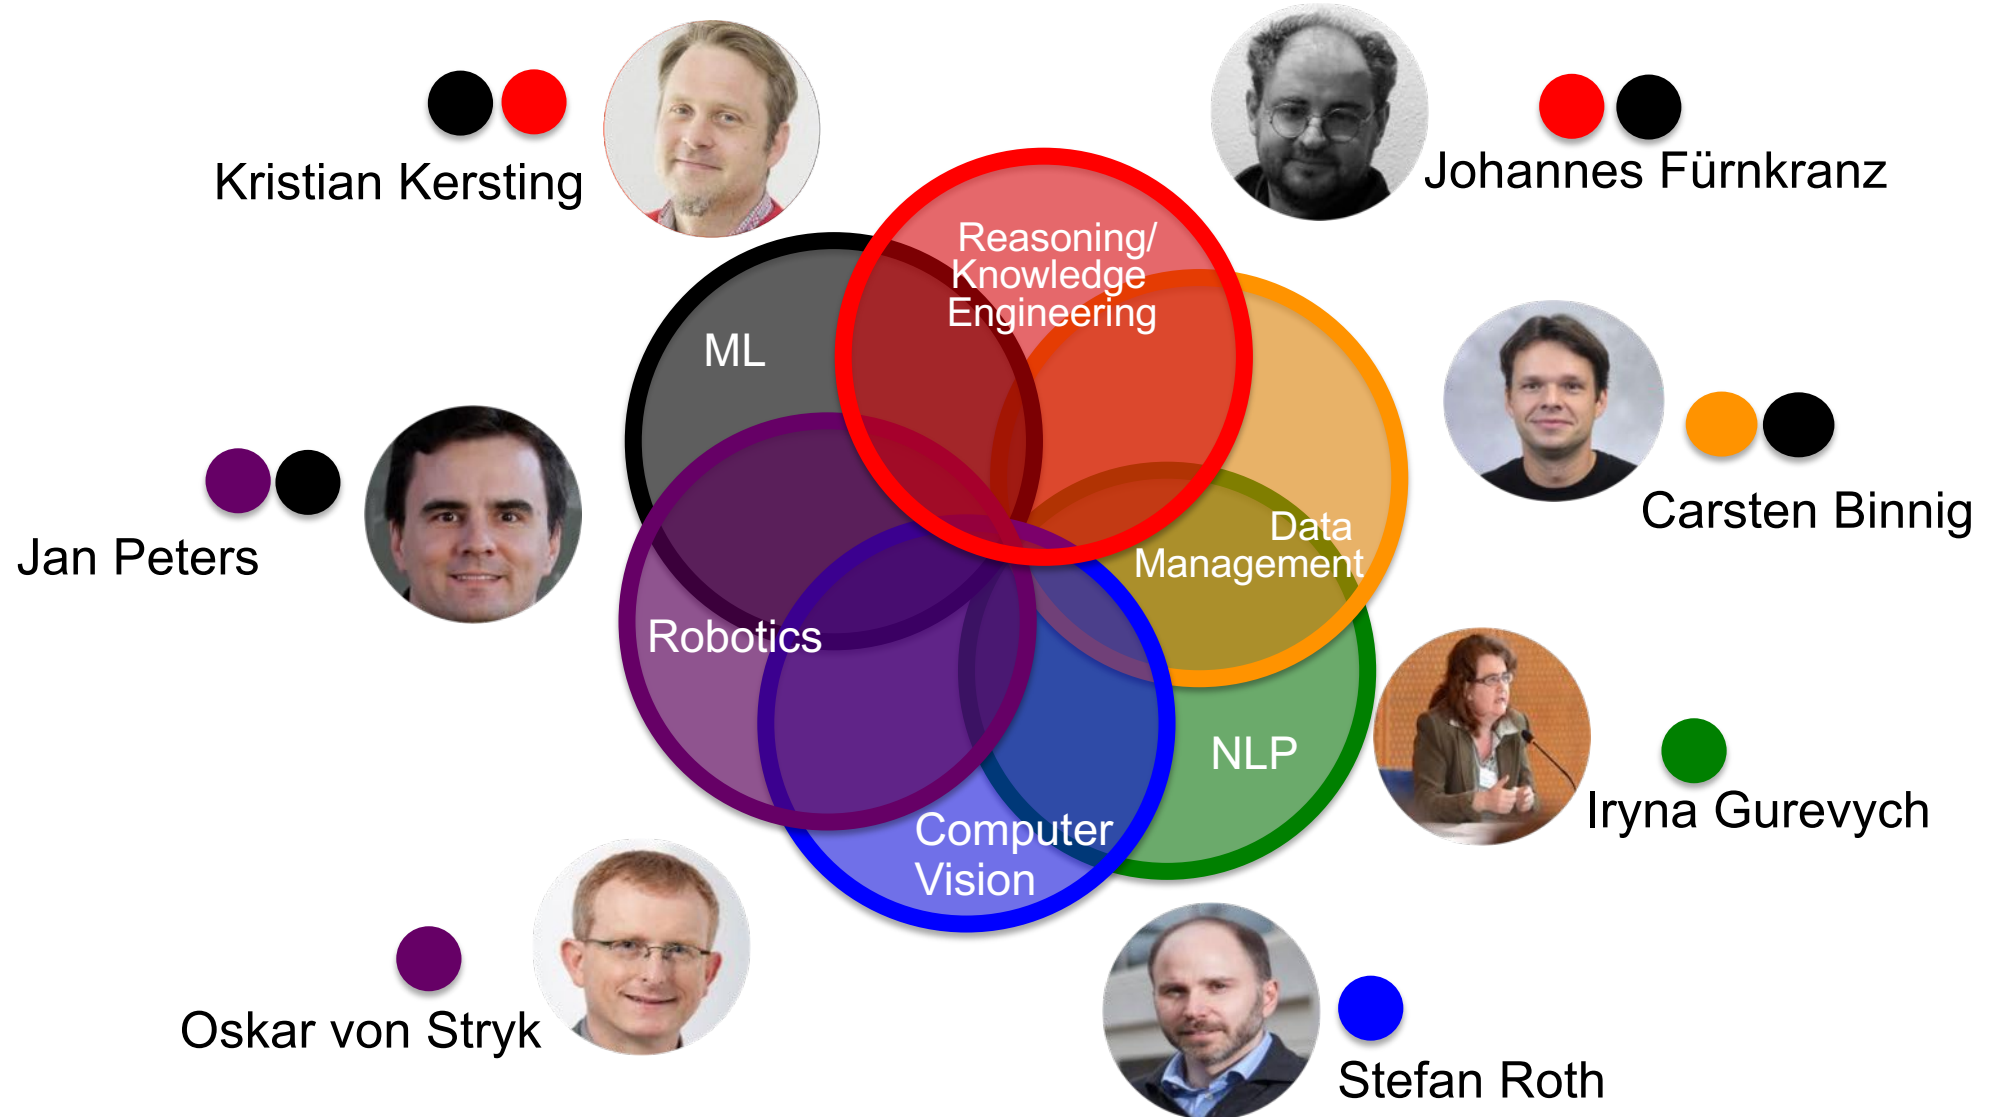
And it appears in industrial strength solvers such as CPLEX and GUROBI



A group of people, likely a sports team, are gathered around a soccer ball. They are wearing blue shirts and are holding the ball with their hands, symbolizing teamwork and collaboration. The background is a green field, possibly a soccer field.

AI is a team sport !

Core AI at TU Darmstadt



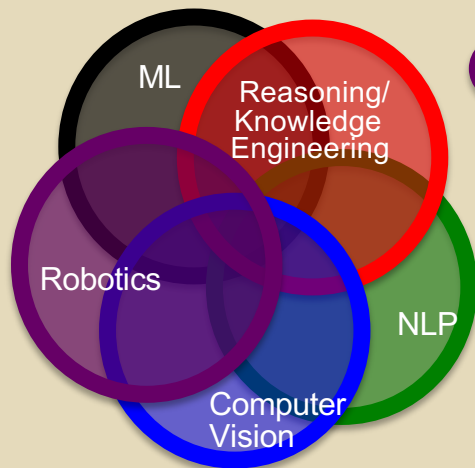
Internationally leading at Statistical Relational AI, Natural Language Processing, Robot Learning, Computer Vision and Machine Learning in general.

Complementary Expertises at TUDA

TUDA Centre for Cognitive Science



“The Machines have to serve the humans.”
AI Strategy of the German government, 15 Nov. 2018



Prof. Dr.
Jan Peters



Prof. Dr.
Stefan Roth



Prof. Dr.
Kristian Kersting

Prof. Dr.
Andre Seyfarth
Locomotion



Prof. Dr.
Heinz Koeppel
Bioinspired Communication Networks

Prof. Dr.
Ralf Galuske
**Systems
Neurophysiology**



Prof. Dr.
Constantin Rothkopf
Psychology of Information Processing



Prof. Dr.
Frank Jaekel
Computational Models of Higher Cognition

Systems AI

The computational and
modeling of complex AI



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

**#1 German and
#2 European
AI institution,
according to csrankings.org**

Meet the AI-DA Scientists

 Machine Learning Relational AI Knuten Kersting	 Robotics Machine Learning Jan Peters	 Visual Inference Stefan Roth
 Systems of AI Data Management Carsten Bittig	 Data Mining Knowledge Engineering Johannes Fürnkranz	 Natural Language Processing Iryna Gurevych
 Computational Models of Human Perception Constantin Rothkopf	 Computational Models of Human Cognition Frank Jeckel	 Systems Optimization Robotics Dekar von Stryk
 Social Human-Robot Interaction Ruth Stock-Homburg	 Software Engineering Mira Meiri	 Embedded Systems Andreas Koch
 Economics and AI Peter Bussmann		

TECHNISCHE
UNIVERSITÄT
DARMSTADT

Centre for
Cognitive
Science

Informations-
Logik
Mathematik