CERN Openlab Machine Learning and Data Analytics WS| April 2016

Machine Learning and Data Analytics at Siemens

Volker Tresp
Siemens Research and Ludwig Maximilian University of Munich

Five Megatrends shaping our world of tomorrow – changes in the markets are accelerating

Demographic change

World population[1], in bn

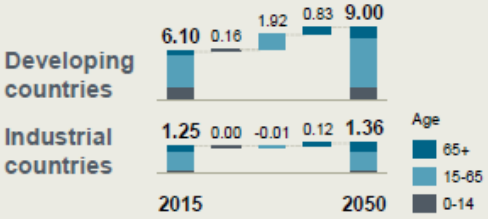

Growing and ageing population

Urbanization

Contribution to global GDP growth, 2007-2025[3], in %

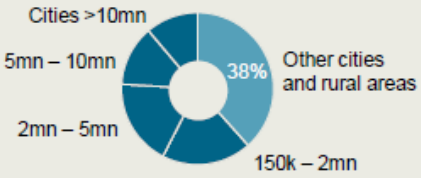

Cities as main driver of GDP growth

Digital transformation

Connected devices[5], in bn

Exponential growth of connected devices ...

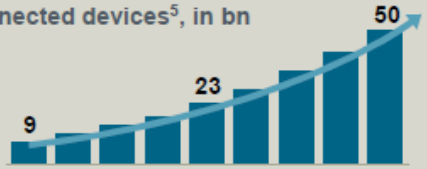

... and digital data[6]

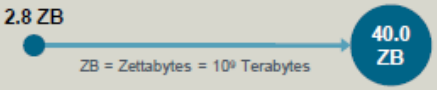

Global warming and weather extremes

Climate change

Annual mean temperature variations 1950-2014[2] (in °C)

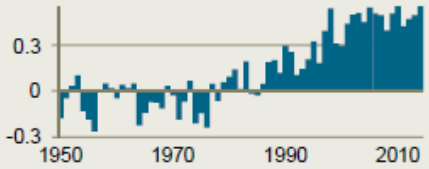

Trend to increase investment abroad

Globalization

Foreign direct investment vs. global GDP, in %[4]

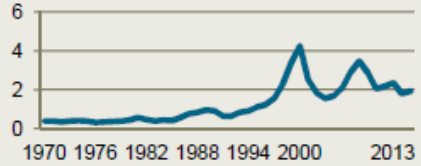

1 UN World Population Prospects (2015)
2 Met Office Hadley Centre observations (2014)
3 McKinsey Global Institute Cityscope (2011)
4 UNCTAD (2013)
5 Cisco: The Internet of Everything (2013)
6 IDC: The Digital Universe (2012)

The Digital Transformation of Services

Classical Time & Material Maintenance → **Performance & Outcome based Contracting** → **Network Platforms**

1. Predictive Maintenance

E.g. Healthcare

Objectives

Prevent unplanned downtime of CT scanners caused by tube failures

2. Performance-based contracts

E.g. Mobility

Objectives

Increase availability and reliability of trains

3. Outcome-oriented contracts

E.g. Power Services

Objectives

Improve customer ROI through flexible service in any market condition

Reactive Maintenance → **Preventive Maintenance** → **Condition-based Maintenance** → **Predictive Maintenance** → **Prescriptive Maintenance**

Sinalytics brings together the technologies needed in an increasingly digitized world

Sinalytics

Data analytics

On-premise, in the cloud and soon in-the-field leveraging Web of Systems technologies

Cyber security

Protecting customer data in open, interconnected industrial IT systems

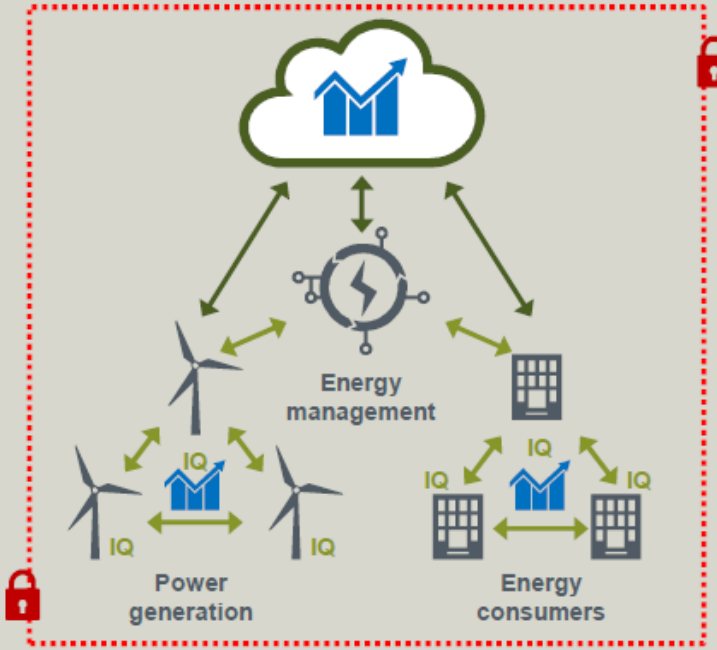

Connectivity

Secure and proven technologies connecting already ~300k devices

Smart networked devices & systems

System-aware, autonomous and app-enabled to meet industry and infrastructure needs

Technology Field: *Business Analytics & Monitoring*

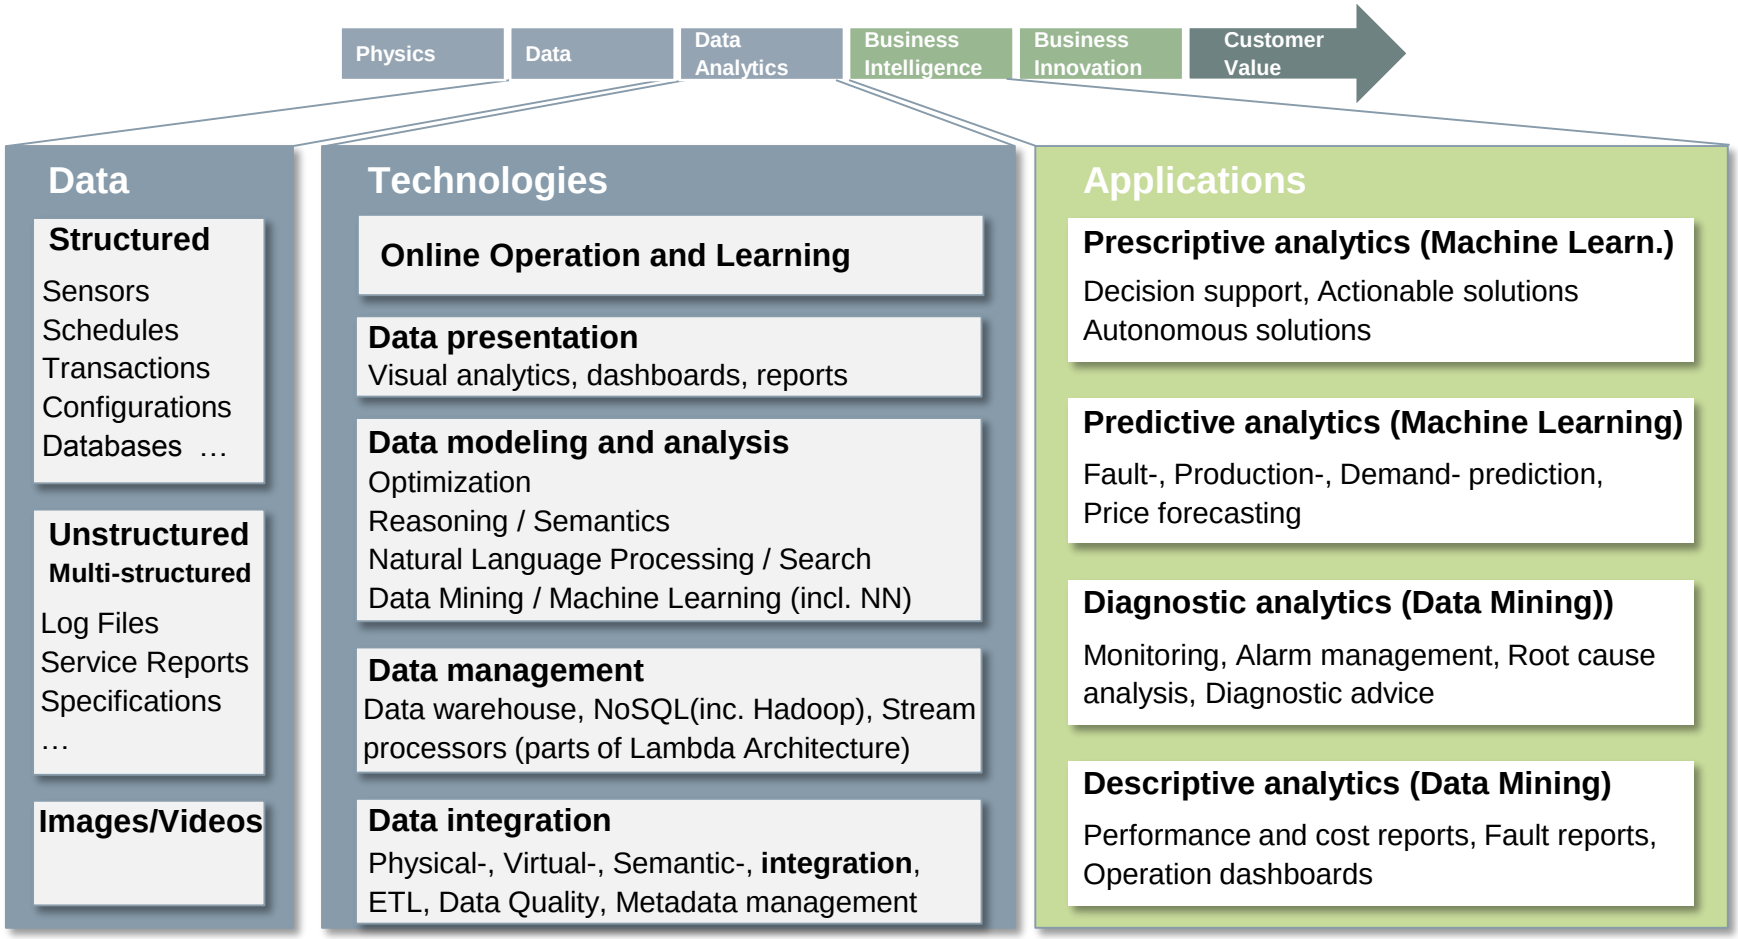

- WatchCat
- Elvis

Two Decades of Experience in the Application of Machine Learning in Challenging Environments

- NeuroSteel
  - “In Field Analytics”

Rolling Mills: no Trivial Task

- How much force *F* is needed for a desired reduction in steel thickness?
  - For all kinds of parameters, and steel properties
- The dependency of the force on 30 and more parameters was modeled with a neural network
- Maybe sounds like a simple problem but:

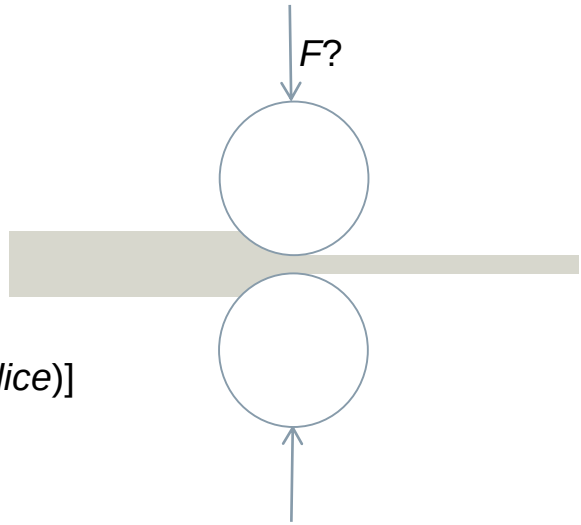

Challenges:

- **Sparse data and safety guarantees**
  - Existing solution was used as a prior
  - Additional data was generated from the existing solution

    [Geoff Hinton referred to the mixing in of prior data as:

    "Priors without Prejudice" (from Jane Austin's *Pride and Prejudice*)]
- **Online learning**
  - The plants change quickly (concept drift)
    - Online learning was required (in general, would be avoided at all costs)
  - Multi-resolution adaptation scheme was developed: performs stable adaptation at different time scales
- **Cold start for a new plant, …**

Röscheisen, Hofmann, Tresp. Neural control for rolling mills: Incorporating domain theories to overcome data deficiency. *NIPS*1991*
Schlang, Feldkeller, Lang, Poppe, and Runkler. Neural Computation in Steel Industry. *European Control Conference (ECC),* 1999

Huge Commercial Success: Examples for World-Wide Installations

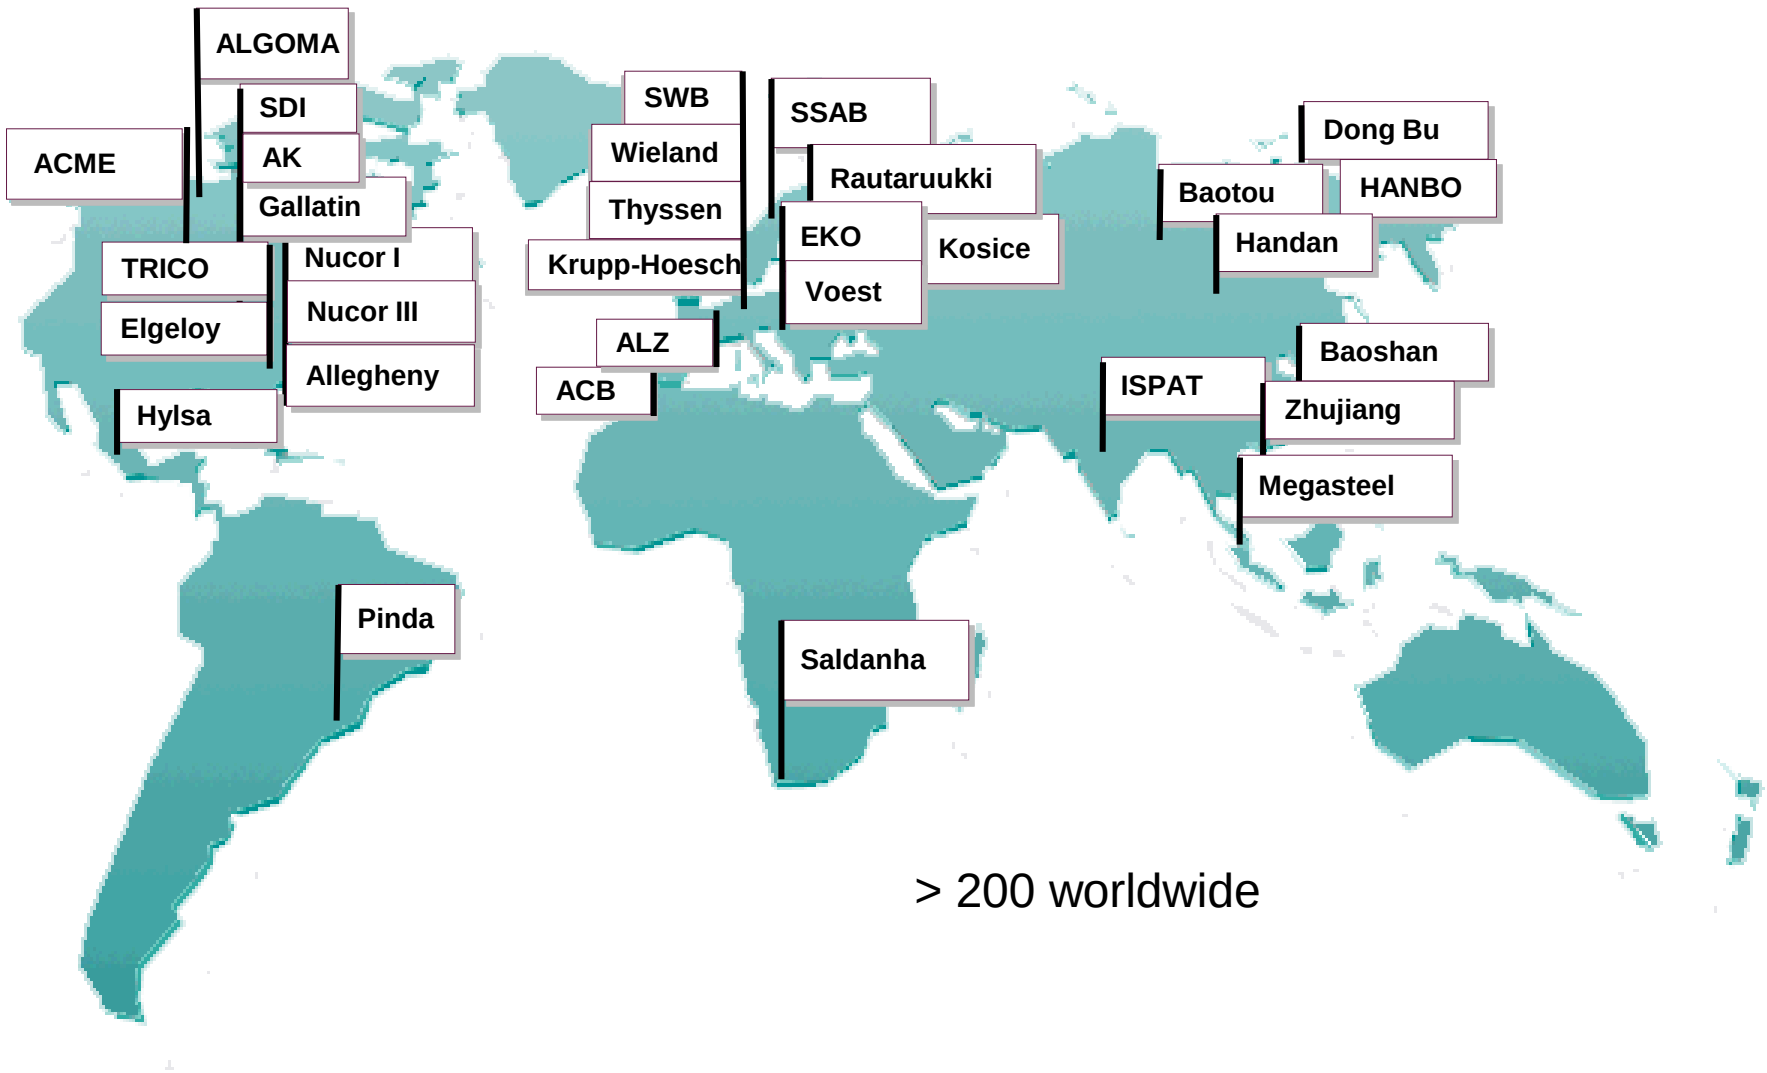

Neural Networks becomes Deep Learning: *First Contact*

- Kai Yu

- Student at LMU (Siemens Stipend) under my supervision ’02
  - Gaussian Processes
- Siemens Research ‘04
- NEC Research ’06:
  - **Deep Neural Networks (in Yann LeCun’s footsteps)**
- BAIDU:
  - **Head of Institute for *Deep Learning*** ‘09
- Horizon Robotics ‘15

Yu Kai, head of Baidu's Institute of Deep Learning (IDL), demonstrates the smart bike project, DuBike, at the company's headquarters in Beijing. Photo: Simon Song

Currently there is Huge Interest in AI in the Public

Baidu's chief scientist explains why computers won't take over the world just yet

by Derrick Harris @derrickharris SEPTEMBER 23, 2015, 11:30 AM EDT

Are AI and "deep learning" the future of, well, everything?

Thanks to the advances in deep machine learning, technology companies across the globe are teaching computers to think for themselves

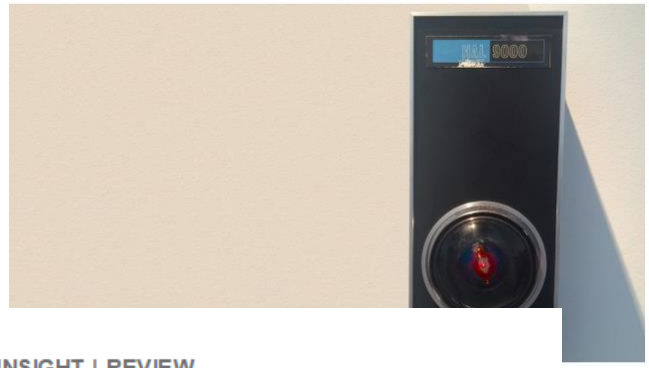

Forbes / Tech

2 FREE Issue

DEC 29, 2014 @ 11:37 AM 82,136 VIEWS

Tech 2015: Deep Learning And Machine Intelligence Will Eat The World

Anthony Wing Kosner

Despite what Stephen Hawking or Elon Musk say, hostile Artificial Intelligence is not going to destroy the world anytime soon. What is certain to happen, however, is the continued ascent of the practical applications of AI, namely deep learning and machine intelligence. The word is spreading in all corners of the tech industry that the biggest part of big data, the unstructured part, possesses learnable patterns that we now have the computing power and algorithmic leverage to discern—and in short order.

For some of life's

NATURE | INSIGHT | REVIEW

Deep learning

Yann LeCun, Yoshua Bengio & Geoffrey Hinton

Affiliations | Corresponding author

Nature **521**, 436–444 (28 May 2015) | doi:10.1038/nature14539
Received 25 February 2015 | Accepted 01 May 2015 | Published

Citation | Reprints | Rights & permissions | Article me

Abstract

Abstract • References • Author information

Deep learning allows computational models that are composed of n

Only for *New Industry*?

- Machine Learning so far had most impact in the new industries: Google, Facebook, Microsoft, ….
- As shown there are a number of applications also in other industries (e.g., at Siemens)

- **“Old industry” only has a future if it embraces Machine Learning**
  - **Autonomous driving**
  - **Digital Health, personalized medicine**
  - **Digitalization (Siemens Business)**
  - **Automation**
  - **Environmental monitoring**

Technology

Deep Learning: Automotive Industry Takes a Giant Leap Forward

Peter Brown

21 July 2015

Share Tweet Email

Recent advancements in artificial neural networks (ANN) and so-called deep learning are accelerating the reality of self-driving vehicles faster than was originally expected as hardware and software vendors are taking the lead in pushing the technology to enable autonomous vehicles forward.

Deep learning, also known as machine learning, has been a concept in place since the early 1980s but only recently has technology advanced to a point where it has become a feasible reality. The idea of deep learning is to attempt tc artificially emulate the functionality of the human brain via hardware and software. An ANN will continuously learn and will base its ability to recognize the surroundings on a deep learning phase based on real examples of sounds, images and input from other senses.

Using deep learning to analyze genetic mutations: an interview with Brendan Frey

Published on September 21, 2015 at 6:07 AM · No Comments

Print PDF Recommend 86 Share 26 Share 4 G+1 Tweet 119

Interview conducted by April Cashin-Garbutt, MA (Cantab)

insights from industry

Brendan Frey

President and CEO of Deep Genomics

Please can you explain what deep learning algorithms are and how they could help to uncover disease-causing genetic mutations?

To understand deep learning in the context of genetic disease, you need to understand shallow learning first. Shallow learning relates mutations to diseases by looking for mutations that commonly occur in patients with a disease. It's a commonly used method.

Using neural network-based forecasting software from Siemens, power generation and demand can be predicted with growing accuracy.

Fanuc Aims to Enhance Factory Robots with ‘Deep Learning’

ARTICLE COMMENTS (1)

ARTIFICIAL INTELLIGENCE DEEP LEARNING FACTORY AUTOMATION FANUC ROBOTICS

Email Print

By TAKASHI MOCHIZUKI

Deep Learning is the Future of Automation and Robotics

Rob Spiegel, Senior Editor, Automation & Motion Control

11/28/2014 2 comments

NO RATINGS
LOGIN TO RATE

Gefällt mir · Tweet 12 Share 7 G+1 6

BIO
EMAIL THIS
PRINT
COMMENT
GOOGLE+

Robots are getting more agile and automation systems are becoming more complex. Yet the most impressive development in robotics and automation is increased intelligence. Machines in automation are increasingly able to analyze huge amounts of data. They are often able to see, speak, even imitate patterns of human thinking. Researchers at the automation company, European Automation, call this deep learning.

Successes are for Real: Computer Vision

- 1000 classes
- 1.28 Mio images
- ***No Feature Engineering!***
- ***An order of magnitude improvement with respect to the state of the art!***

Model	Top-1	Top-5
Sparse coding [2]	*47.1%*	*28.2%*
SIFT + FVs [24]	*45.7%*	*25.7%*
CNN	**37.5%**	**17.0%**

Table 1: Comparison of results on ILSVRC-2010 test set. In *italics* are best results achieved by others.

- Human: 5%
- "Sensational" 33% improvement by Alex-Net in 2010
- 2015: 86% improvement!

MSR (Dec 2015):

considerably increased depth. On the ImageNet dataset we evaluate residual nets with a depth of up to 152 layers—8× deeper than VGG nets [41] but still having lower complexity. An ensemble of these residual nets achieves 3.57% error on the ImageNet test *set. This result won the 1st place on the ILSVRC 2015 classification task. We also present analysis*

Deep Learning comes with an incredible powerful infrastructure

- Theano, Keras, Caffe, Torch , TensorFlow… are powerful software development infrastructures
- GPU computing can be used to speed up computation
- FPGA … can be used for implementing trained networks
- The community itself:
  - Explosion of ideas and creativity
  - Scientists and engineers and ….

Decaf / Caffe

a Berkeley Vision Project

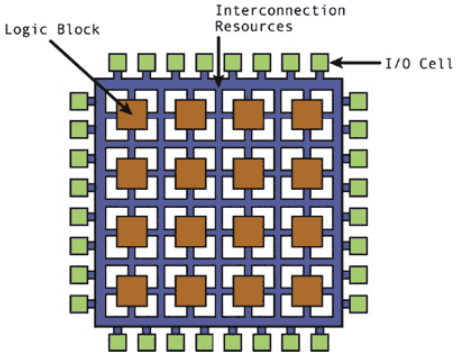

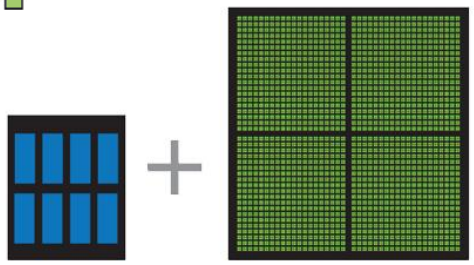

DL@Siemens (the Future)



- Setting up a powerful computing infrastructure
- Weekly Journal Club
- Demonstrators
- Impact in Siemens

"Innovating today is about creativity, it is about the freedom to act…if you look at the big conglomerate…it's highly regulated, it's a lot of talk about…internal controls and this and that…so there's a huge mindshift change to get the best of people" **Kaeser**

Siemens Activities: Not Starting with Zero

Other ongoing activities at Siemens Research (excerpt):

- Time-series modeling and prediction with recurrent (deep) neural networks with a 25 year history (Zimmermann, Grothmann)
- Optimization and condition monitoring of gas- and wind turbines with reinforcement learning and recurrent neural networks (Sterzing, Udluft, Hentschel, Tokic)
- SENN Framework

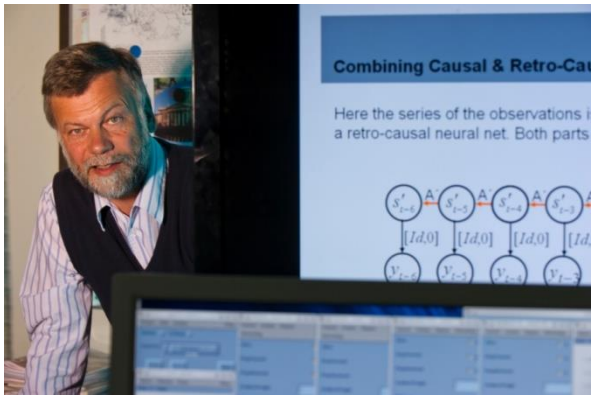

- Medical Image Analysis at Siemens Healthcare (Comaniciu)

Reinforcement Learning

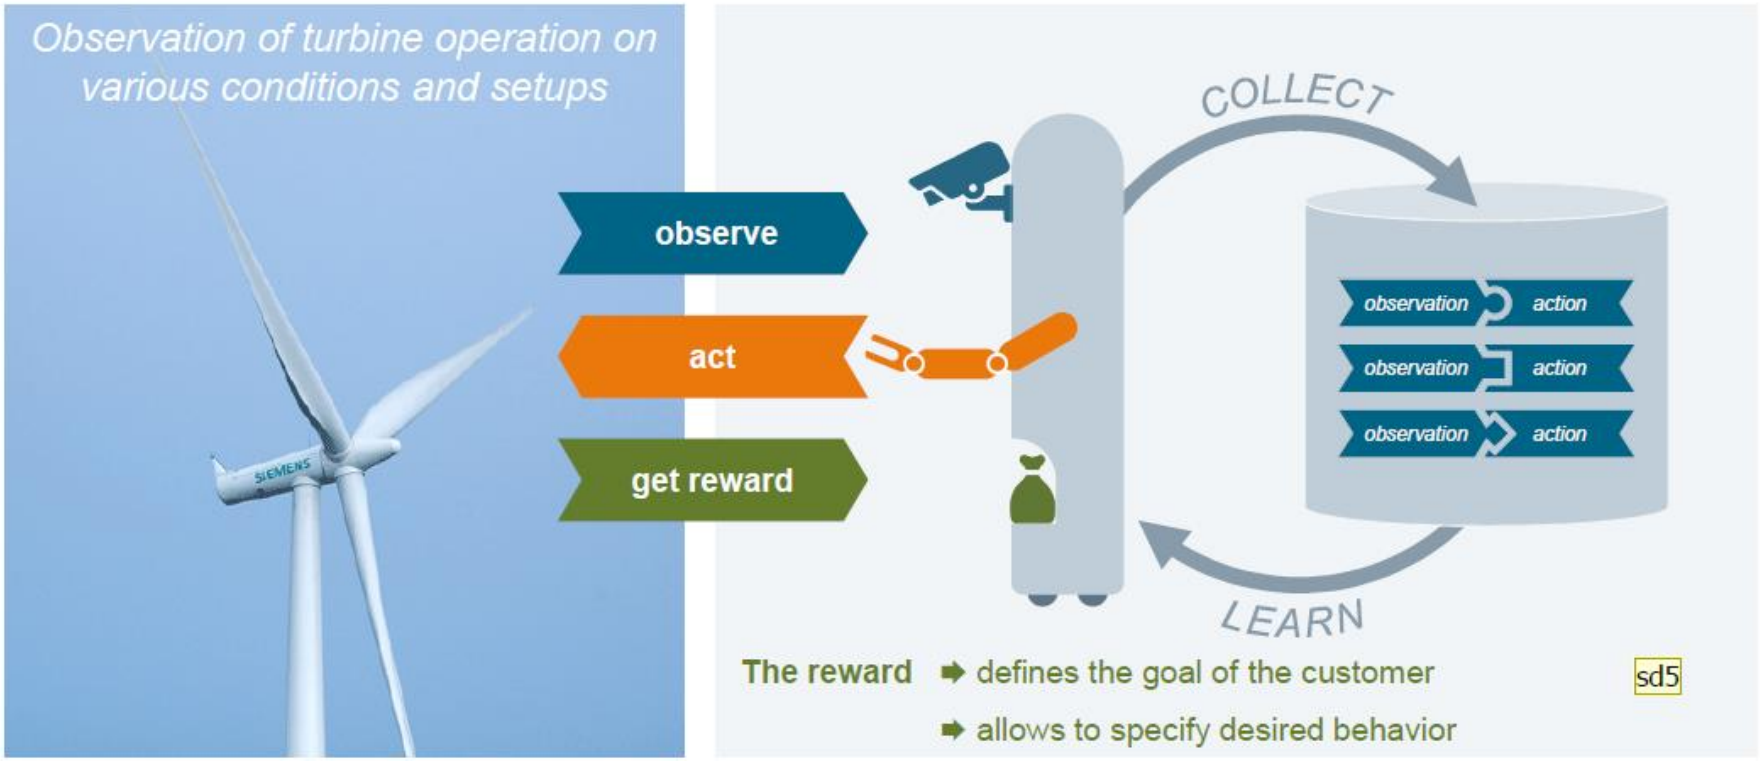

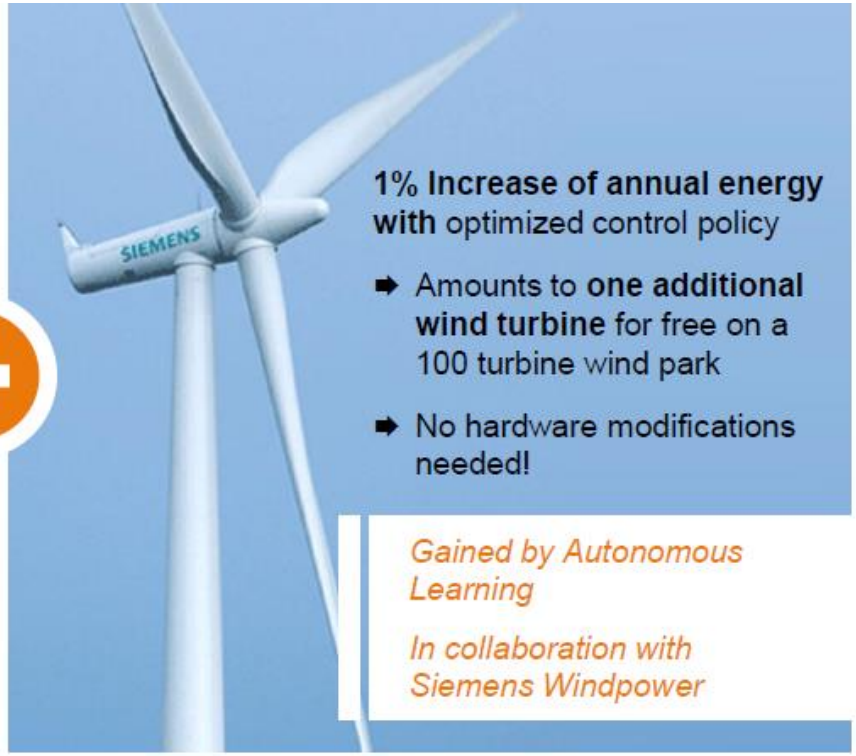
SIEMENS
1% Increase of annual energy with optimized control policy
Amounts to one additional wind turbine for free on a 100 turbine wind park
No hardware modifications needed!
Gained by Autonomous Learning
In collaboration with Siemens Windpower

With a Little Help from our Brain

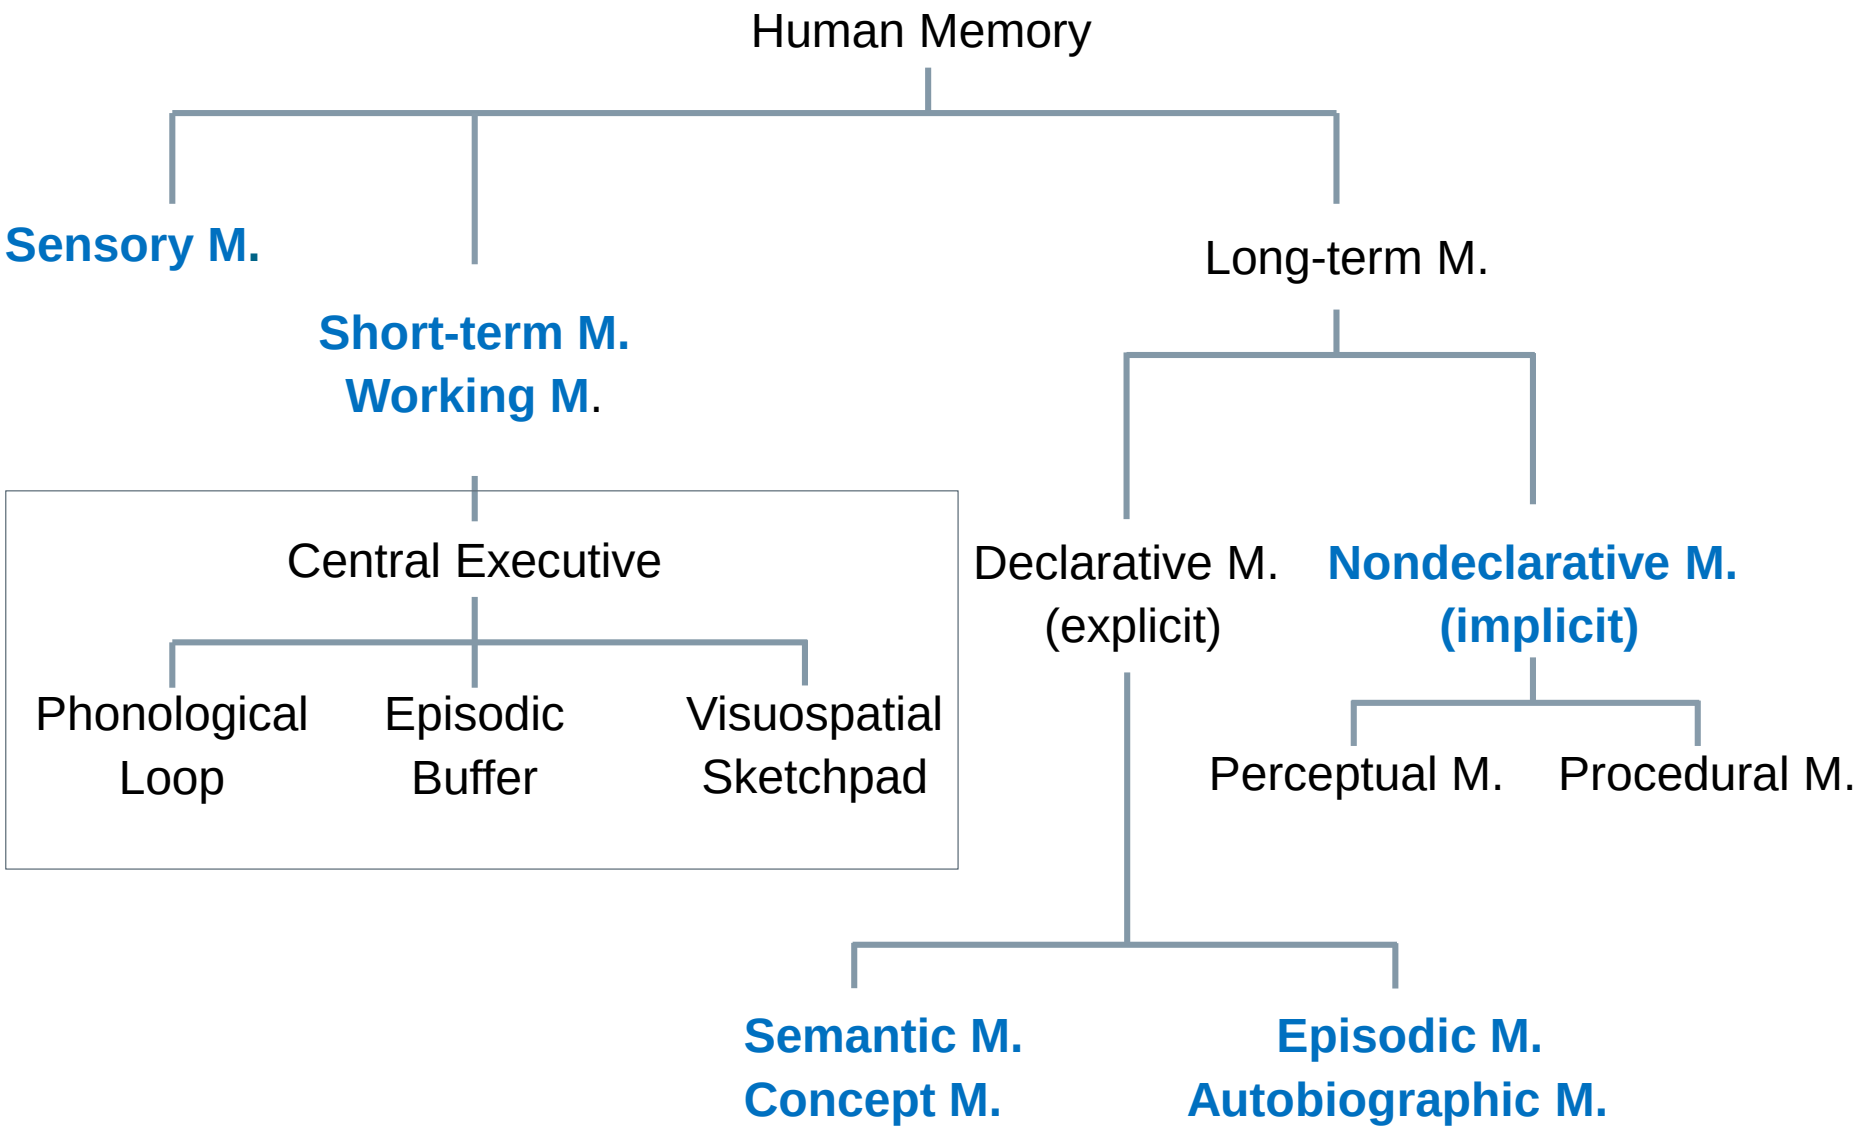
Human Memory
Sensory M.
Short-term M.
Working M.
Long-term M.
Central Executive
Phonological Loop
Episodic Buffer
Visuospatial Sketchpad
Declarative M. (explicit)
Nondeclarative M. (implicit)
Perceptual M.
Procedural M.
Semantic M.
Concept M.
Episodic M.
Autobiographic M.

Deep Learning Ecosystem for High-Dimensional (Sparse) Data

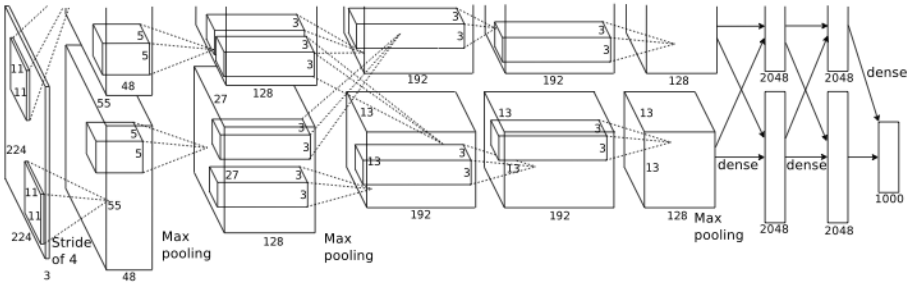

Deep Convolutional Networks exploit locality in time and space and combine local features to form flexible complex patterns

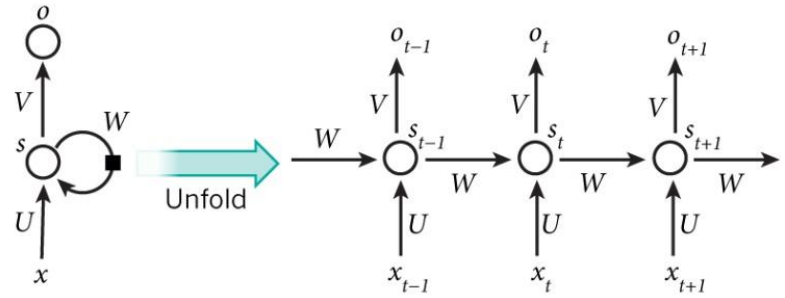

Recurrent Networks exploit locality in time and space and combine local (short term) features to form complex (long term) patterns by exploiting internal memory

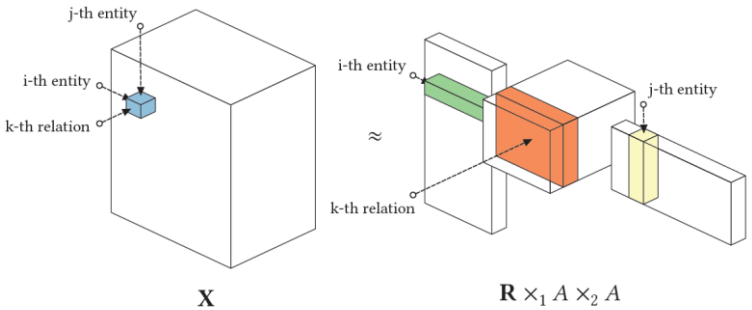

$$\mathbf{R} \times_1 A \times_2 A$$

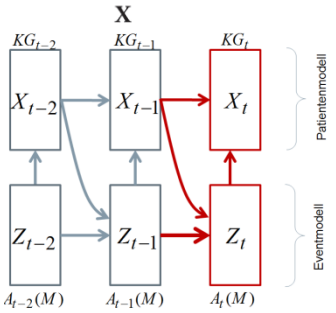

Multiway Neural Networks

Tensor Decompositions for

- **Sensory Memory Models**
- **Semantic Knowledge Graph Models**
- **Episodic Event Memories**

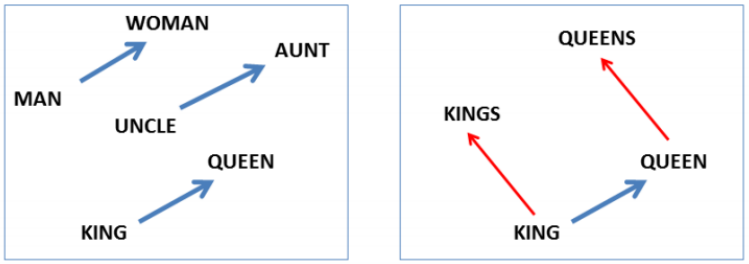

(Mikolov et al., NAACL HLT, 2013)

Representation Learning to achieve latent representations of words, entities and events

External data sources

Internal data sources

Siemens Corporate Knowledge Graph

Knowledge graph set up in secure infrastructure

Smart Data Lab

SDW dev server

SDW Queries

Query	Show me all **suppliers** of **optical sensors**
Query	What is the **parent company** of supplier X
Query	Which **sensor suppliers** have an **order volume** larger than 100k EUR

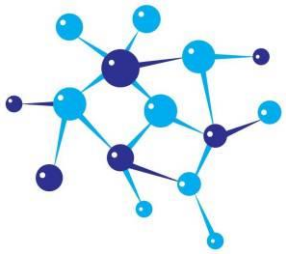

Clinical Data Intelligence

Klinische Datenintelligenz

Funded by the
Federal Ministry for Economic Affairs and Energy

Technologiy Programm “Smart Data”

Bundesministerium für Wirtschaft und Energie

Smart Data

SIEMENS

CHARITÉ
UNIVERSITÄTSMEDIZIN BERLIN

Universitätsklinikum Erlangen

FAU FRIEDRICH-ALEXANDER UNIVERSITÄT ERLANGEN-NÜRNBERG

Fraunhofer IIS

DFKI Deutsches Forschungszentrum für Künstliche Intelligenz GmbH

averbis
medical language technology

Institut für FrauenGesundheit (IFG®)

Dynamic Multiway Neural Network for Medical Decision Processes (Nephrology)

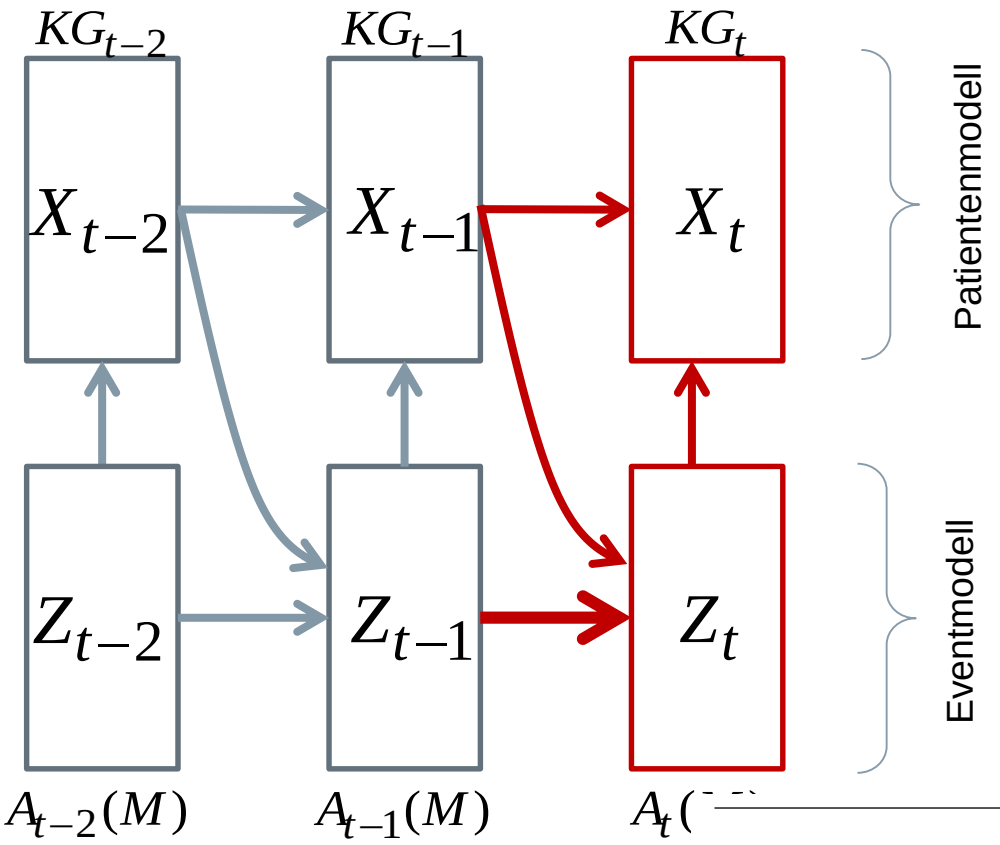

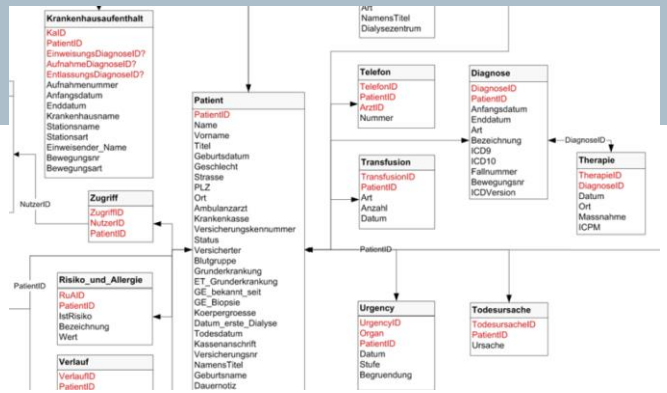

	AUPRC	AUROC
Markov-tensor	0.574 ± 0.0014	0.977 ± 0.0001
Logistic Regression	0.554 ± 0.0020	0.970 ± 0.0005
KNN	0.482 ± 0.0012	0.951 ± 0.0002
Naive Bayes	0.432 ± 0.0019	0.843 ± 0.0015
Constant predictions	0.350 ± 0.0011	0.964 ± 0.0001
Random	0.011 ± 0.0001	0.5

- Analysis
- Prediction
- Prescription

	AUPRC	AUROC
KG+Markov	0.586 ± 0.0010	0.979 ± 0.0001
Markov	0.574 ± 0.0014	0.977 ± 0.0001
KG	0.487 ± 0.0016	0.974 ± 0.0002

Esteban, Schmidt, Krompaß, Tresp. Predicting Sequences of Clinical Events by using a Personalized Temporal Latent Embedding Model. *IEEE ICHI*, 2015

Conclusions

- There is a long history of machine learning and data analytics at Siemens
- Digitization is the main business driver of innovation
- The Deep Learning Ecosystem is a main technological driver of innovation: CNN, RNN, multiway NN, Representation Learning
- We are looking forward to continuing the fruitful collaboration with Industrial Control Systems team

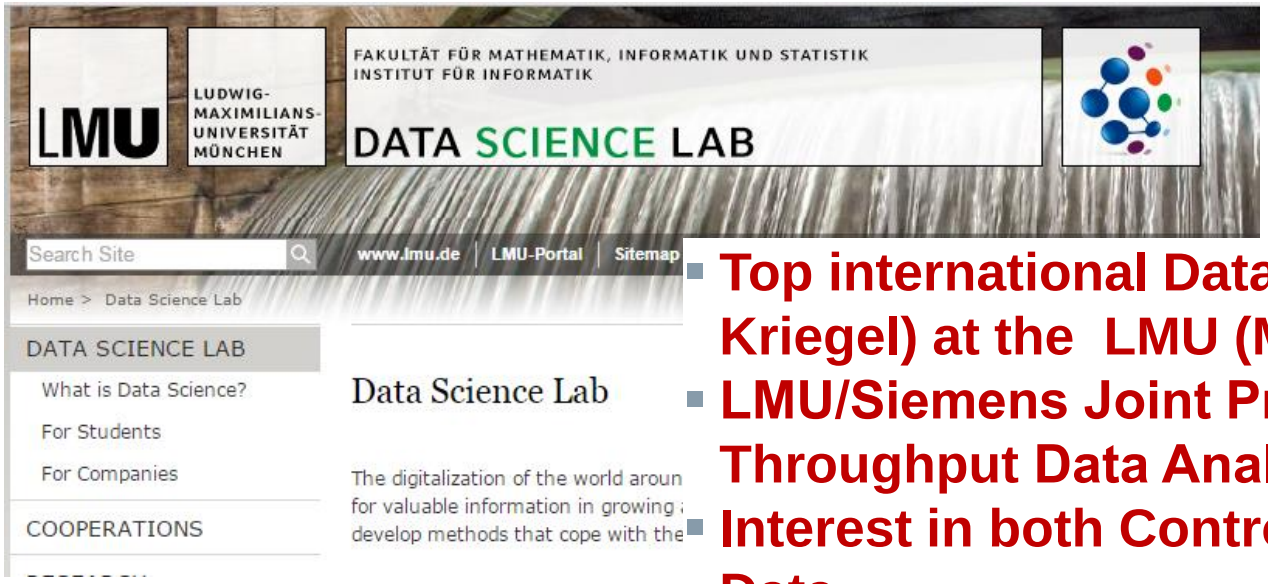

- **Top international Datamining Team (H.P. Kriegel) at the LMU (Munich Univ.)**
- **LMU/Siemens Joint Project on „High Throughput Data Analysis“**
- **Interest in both Control Data and Physics Data**